Monitoring Your CMS Tier 3 Site

Joel W. Walker

Sam Houston State University

OSG and CMS Tier 3 Summer Workshop

Texas Tech University

August 9-11, 2011

Representing:
the Texas A&M Tier 3 CMS Grid Site on the Brazos Cluster

In Collaboration With:
David Toback
Guy Almes
Steve Johnson
Jacob Hill
Michael Kowalczyk
Vaikunth Thukral (With thanks for marked slides)
Daniel Cruz

Introduction to Grid Computing

- Cluster
  - Multiple computers in a Local Network
- The Grid
  - Many clusters connected by a Wide Area Network
  - Resources expanded for thousands of users as they have more access to distributed computing and disk
- CMS Grid: Tiered Structure (Mostly about size & location)
  - Tier 0: CERN
  - Tier 1: A few National Labs
  - Tier 2: Bigger University Installations for national use
  - Tier 3: For local use (Our type of center)

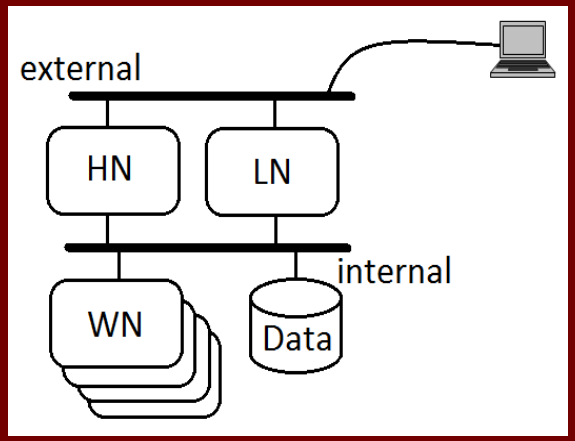

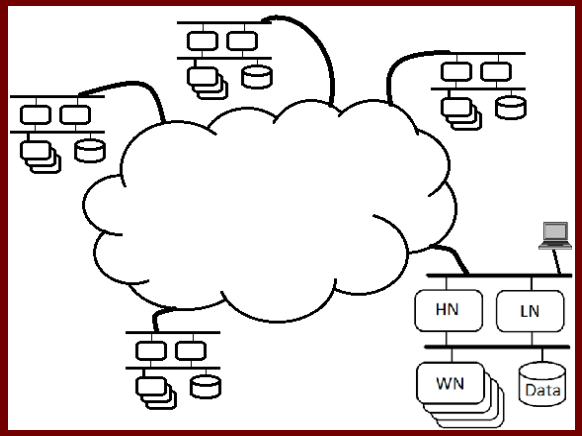

Advantages of Having a CMS Tier 3 Computing Center at TAMU

- Don’t have to compete for resources
  - CPU priority - Even though we only bought a small amount of CPUs, can periodically run on many more CPUs at the cluster at once
  - Disk space - Can control what data is here
- With a “standardized” Tier 3 on a cluster, can run same here as everywhere else
- Physicists don’t do System Administration

T3_US_TAMU as part of Brazos

- Brazos cluster already established at Texas A&M
- Added our own CMS Grid Computing Center within the cluster
- Named T3_US_TAMU as per CMS conventions

T3_US_TAMU added CPU and Disk to Brazos as our way of joining

- Disk
  - Brazos has a total of ~150TB of storage space
  - ~30 TB is assigned to our group
  - Space is shared amongst group members
    - N.B. Another 20TB in the works
- CPU
  - Brazos has a total of 307 compute nodes/2656 cores
  - 32 nodes/256 cores added by T3_US_TAMU
    - Since we can run 1 job on each core → 256 jobs at any one time, more when cluster is underutilized, or by prior agreement
  - 184,320 (256 x 24 x 30) dedicated CPU hours/Month

Motivation

1) Every Tier 3 site is a unique entity composed of a vast array of extremely complicated interdependent hardware and software, extensively cross-networked for participation in the global endeavor of processing LHC data.

Motivation

2) Successful operation of a Tier 3 site, including performance optimization and tuning, requires intimately detailed, near real-time feedback on how the individual system components are behaving at a given moment, and how this compares to design goals and historical norms.

Motivation

3) Excellent analysis tools exist for reporting most of the crucial information, but they are spread across a variety of separate pages, and are designed for generality rather than site-specificity. The quantity of information can be daunting to users, and not all of it is useful. A large amount of time is spent clicking, selecting menus, and waiting for results, and it is still difficult to be confident that you have obtained the "big picture" view.

Funding

• The TAMU Tier 3 Monitoring project is funded by a portion of the same grant which was used to purchase the initial "buy in" servers added to the Brazos cluster. It represents an exciting larger school – smaller school collaboration between Texas A&M and Sam Houston State University.

• The funding represents a generous one time grant by the Norman Hackermann Advanced Research Project, an internally awarded entity of the Texas Higher Education Coordinating Board. (They love big-small collaborations!)

• The Co-PI's are Dr. Dave Toback (Physics) and Dr. Guy Almes (Computer Science), both of Texas A&M University

Monitor Design Philosophy and Goals

1) The monitor must consolidate all key metrics into a single clearing house, specialized for the evaluation of a single Tier 3 site.

Monitor Design Philosophy and Goals

2) The monitor must provide an instant visually accessible answer to the primary question of the operational status of key systems.

Monitor Design Philosophy and Goals

3) The monitor must facilitate substantial depth of detail in the reporting of system behavior when desired, but without cluttering casual usage, and while providing extremely high information density.

Monitor Design Philosophy and Goals

4) The monitor must provide near real-time results, but should serve client requests immediately, without any processing delay, and without the need for the user to make parameter input selections.

Monitor Design Philosophy and Goals

5) The monitor must allow for the historical comparison of performance across various time scales, such as hours, days, weeks, and months.

Monitor Design Philosophy and Goals

6) The monitor must proactively alert administrators of anomalous behavior.

… This is currently the only design goal which still lacks at least a partial implementation. The others are at least "nearly done".

How Does it Work?

A team of CRON – activated Perl scripts harvest the relevant data and images from the web at regulary intervals (currently every 30 minutes, except for the longer interval plots). Most required pages are accessible via CMSWeb.Cern.Ch (PhEDEx Central, and the CMS Dashboard Historical View, Task Monitoring, Site Availability, and Site Status Board), but we also query custom cgi-bin scripts hosted at Brazos.Tamu.Edu for the local execution of "qstat" and "du".

How Does it Work?

These scripts store retrieved images locally for rapid redeployment, including resized thumbnails which are generated “on the fly”. They also compile and sort the relevant information needed to create custom table format summaries, and write the html to static files which will be “included” (SSI) into the page served to the client. The data combined into a single custom table may in some cases represent dozens of recursively fetched webpages.

Is There a Demonstration Version Accessible?

The "Brazos Tier 3 Data Transfer and Job Monitoring Utility" is functioning, although still under development, and the current implementation is openly accessible on the web:

collider.physics.tamu.edu/tier3/mon/

Please open up a web browser and follow along!

Division of Principal Monitoring Tasks

- **I - Data Transfers to the Local Cluster**

  … PhEDEx Transfer Rate and Quality

- **II - Data Holdings on the Local Cluster**

  … PhEDEx Resident and Subscribed Data, plus the local unix “du” reports

- **III - Job Status of the Local Cluster**

  … net job count, CRAB tests, SAM heuristics, CPU usage, and termination status summaries

I - Data Transfers to the Local Cluster

PhEDEx

- **Ph**ysics **E**xperiment **D**ata **Ex**port
  - Data is spread around the world
  - Can Transport tens of Terabytes of data to A&M per month

PhEDEx at Brazos

- PhEDEx performance is continually tested in different ways:
  - LoadTests
  - Transfer Quality
  - Transfer Rate

Production Data	Load Tests
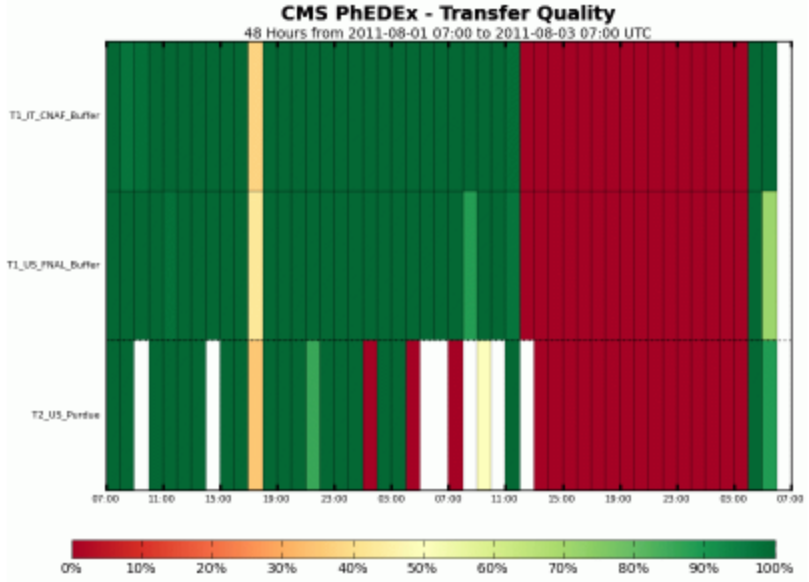	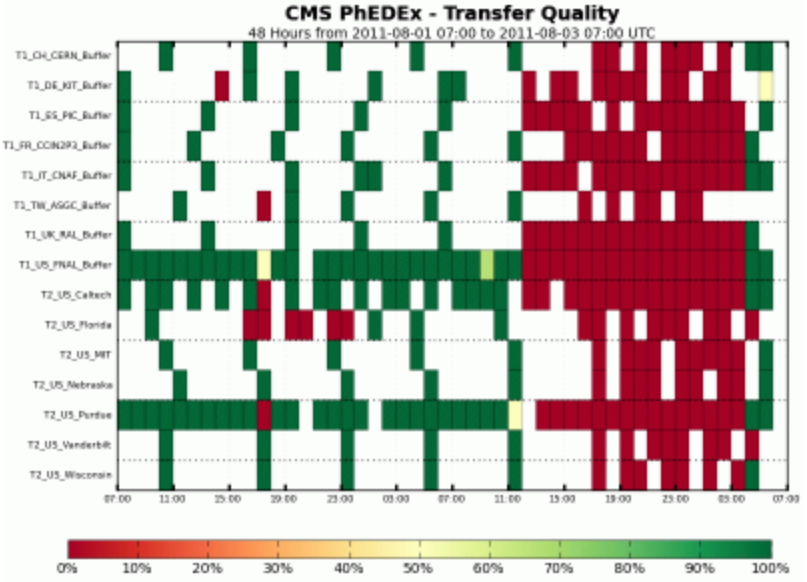
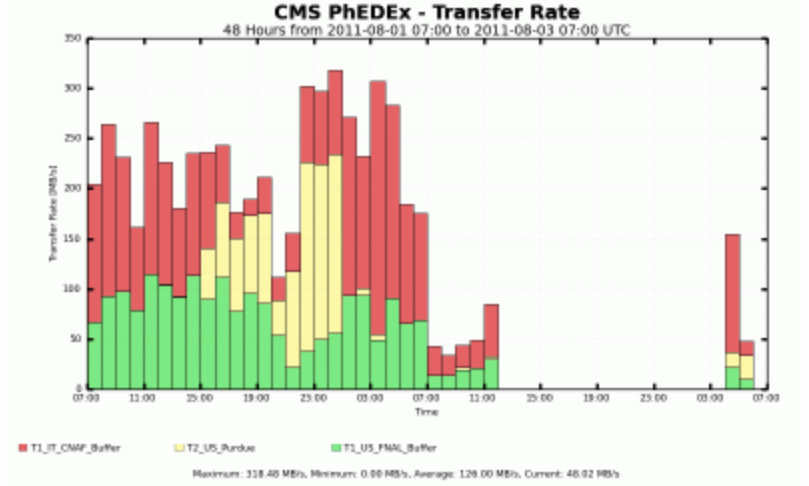	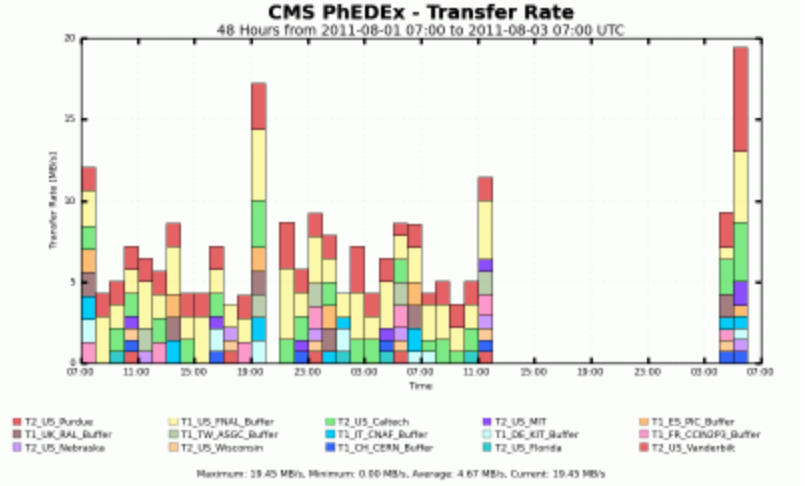
↑ Click to Enlarge Images	Select → Hourly Daily Weekly

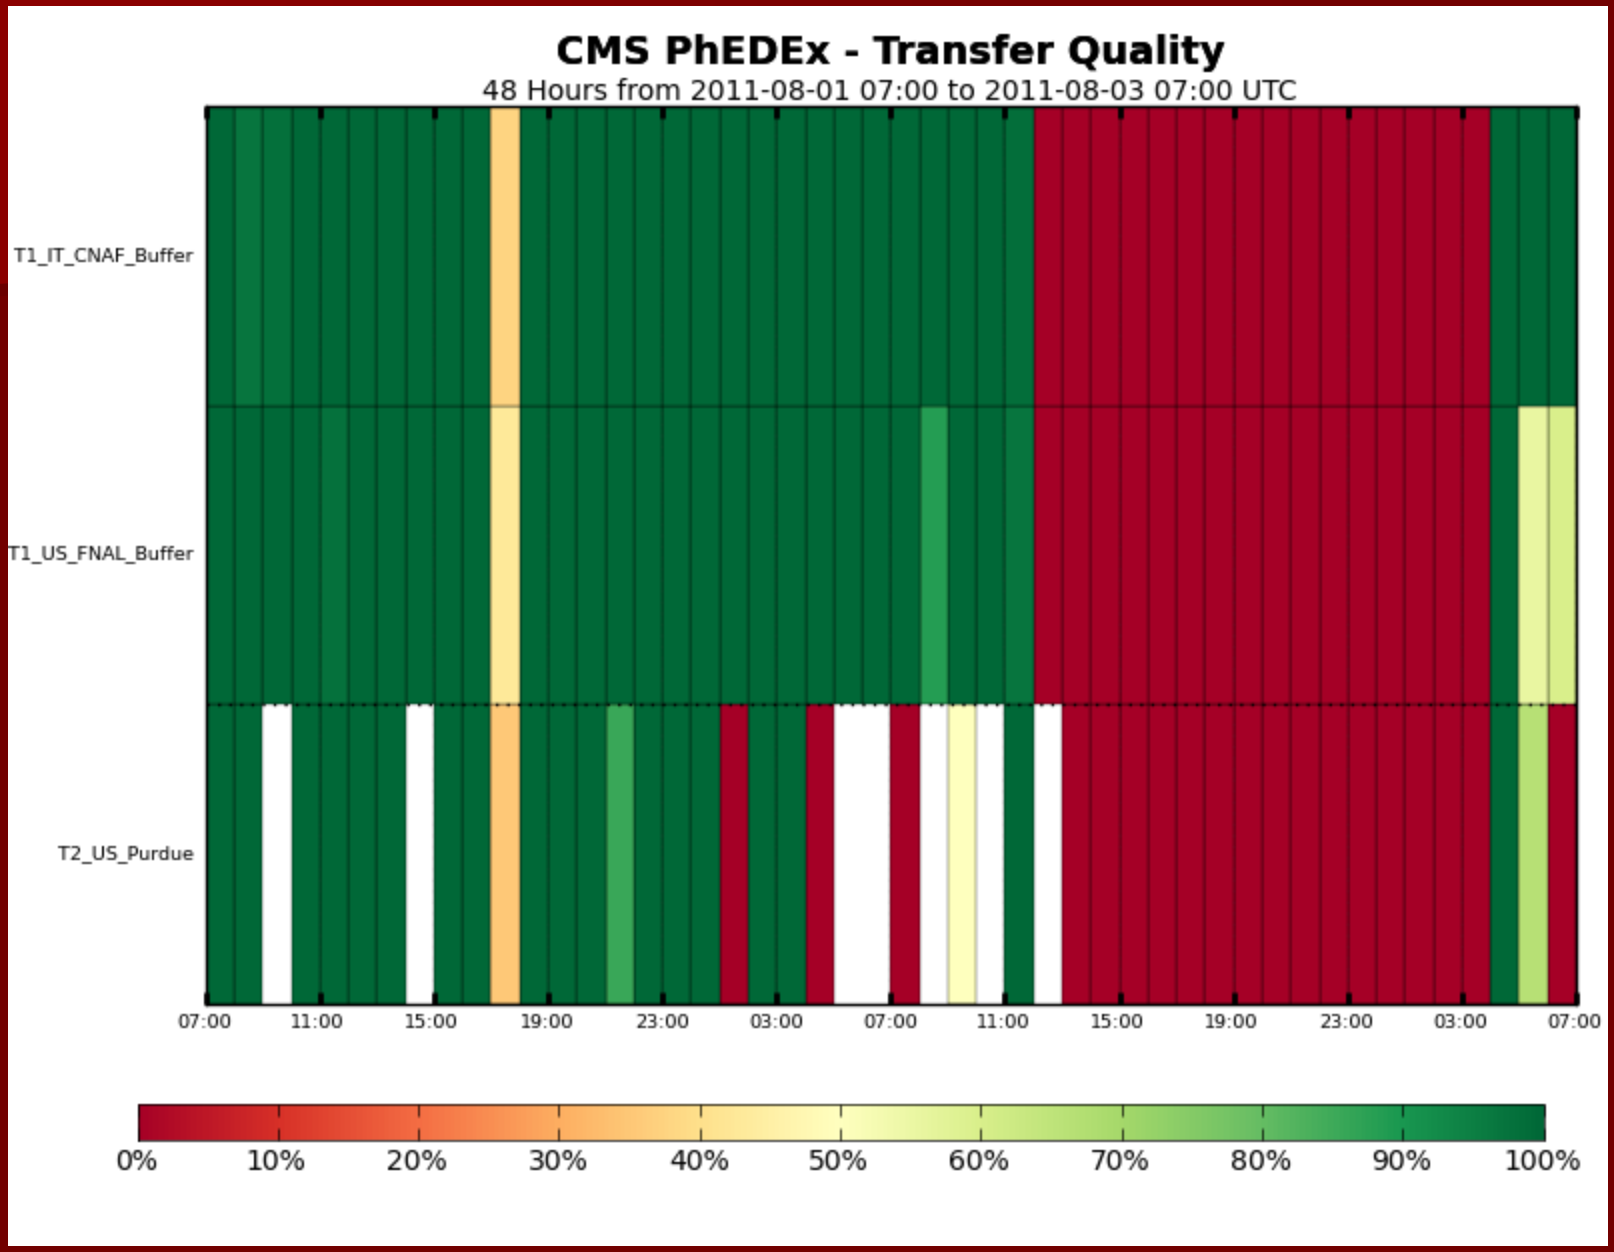
CMS PhEDEx - Transfer Quality
48 Hours from 2011-08-01 07:00 to 2011-08-03 07:00 UTC
T1_IT_CNAF_Buffer
T1_US_FNAL_Buffer
T2_US_Purdue
07:00
11:00
15:00
19:00
23:00
03:00
07:00
11:00
15:00
19:00
23:00
03:00
07:00
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%

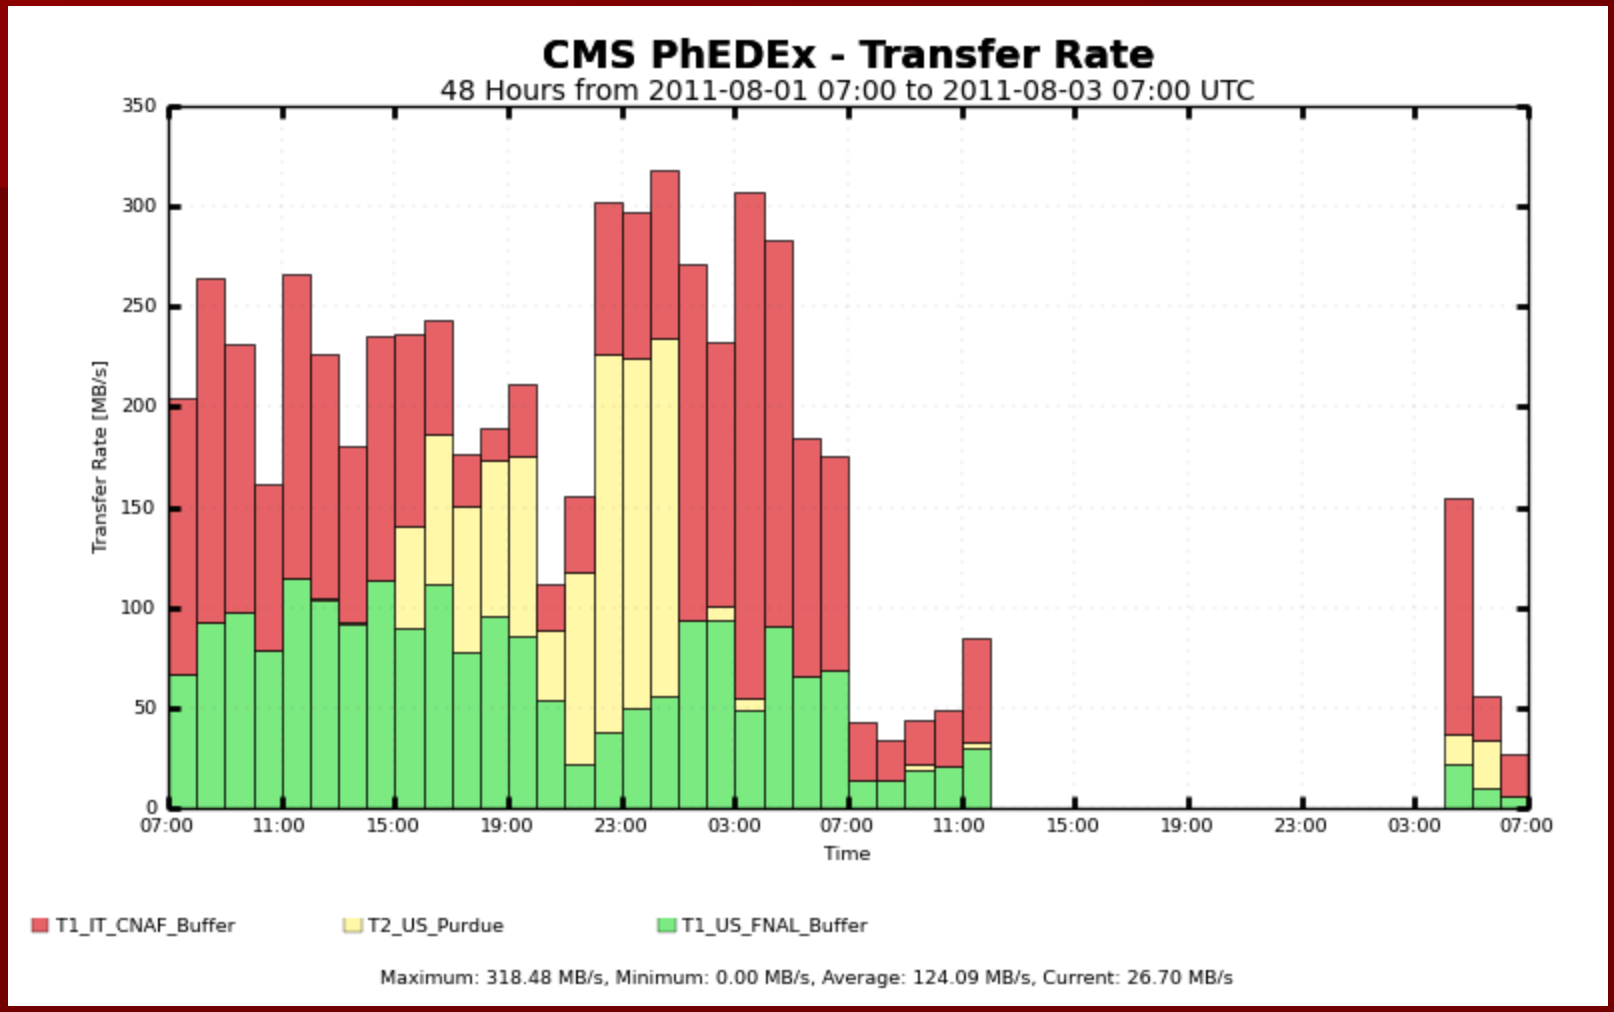
CMS PhEDEx - Transfer Rate
48 Hours from 2011-08-01 07:00 to 2011-08-03 07:00 UTC
Transfer Rate [MB/s]
350
300
250
200
150
100
50
0
07:00
11:00
15:00
19:00
23:00
03:00
07:00
11:00
15:00
19:00
23:00
03:00
07:00
Time
T1_IT_CNAF_Buffer
T2_US_Purdue
T1_US_FNAL_Buffer
Maximum: 318.48 MB/s, Minimum: 0.00 MB/s, Average: 124.09 MB/s, Current: 26.70 MB/s

↓ Production	PhEDEx Data Transfers	Load Test ↓	← Select →	Hour	Day	Week	Month
Link Status	Linked Node	Link Status	Rate	Bytes	Files	Expired	Errors
Valid	T1_CH_CERN_Buffer	Valid	-	-	-	-	-
Valid	T1_DE_KIT_Buffer	Valid	-	-	-	-	-
Valid	T1_ES_PIC_Buffer	Valid	-	-	-	-	-
Valid	T1_FR_CCIN2P3_Buffer	Valid	-	-	-	-	-
Valid	T1_IT_CNAF_Buffer	Valid	15.9 MiB/s	1.3 TiB	577	638	1834
Valid	T1_TW_ASGC_Buffer	Valid	-	-	-	-	-
Valid	T1_UK_RAL_Buffer	Valid	-	-	-	-	-
Valid	T1_US_FNAL_Buffer	Valid	8.2 MiB/s	688.1 GiB	459	0	311
Valid	T2_IT_Pisa	Valid	-	-	-	-	-
Valid	T2_US_Caltech	Valid	-	-	-	-	-
Valid	T2_US_Florida	Valid	-	-	-	-	-
Valid	T2_US_MIT	Valid	-	-	-	-	-
Valid	T2_US_Nebraska	Valid	-	-	-	-	-
Valid	T2_US_Purdue	Valid	1.9 MiB/s	159.5 GiB	16	0	175
Valid	T2_US_UCSD	Valid	-	-	-	-	-
Valid	T2_US_Vanderbilt	Valid	-	-	-	-	-
Valid	T2_US_Wisconsin	Valid	-	-	-	-	-
Valid	T3_US_Colorado	Valid	-	-	-	-	-
Valid	T3_US_Rice	Valid	-	-	-	-	-
Valid	T3_US_TTU	Excluded	-	-	-	-	-
Totals			26.0 MiB/s	2.1 TiB	1052	638	2320

II - Data Holdings on the Local Cluster

Production Data	Load Tests
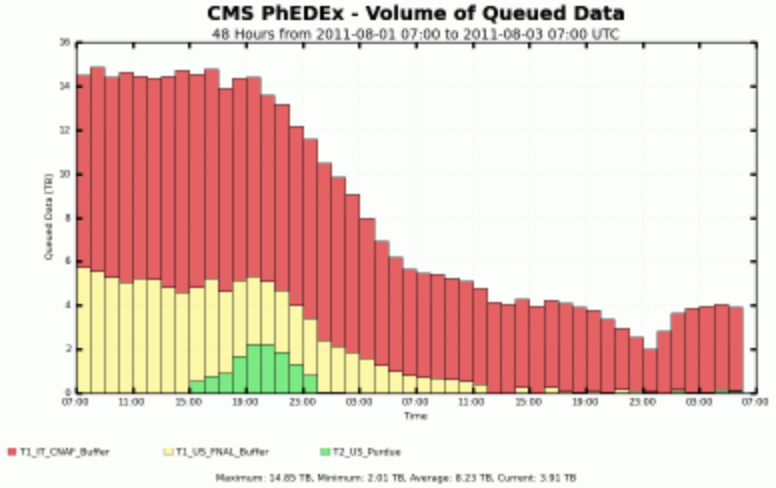	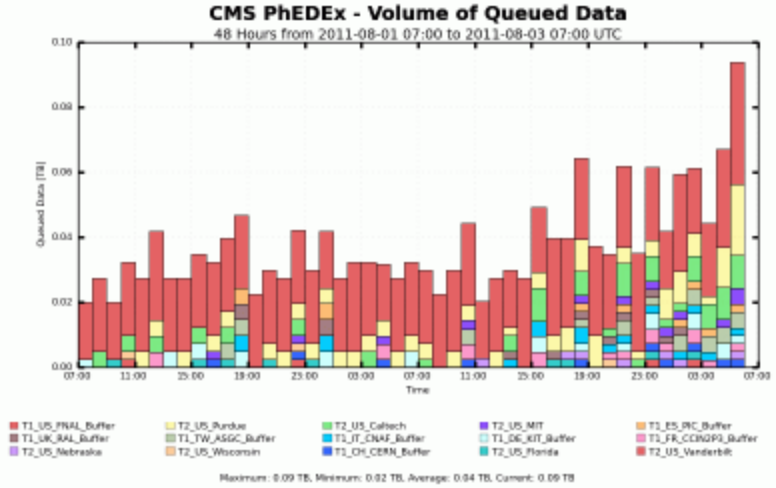
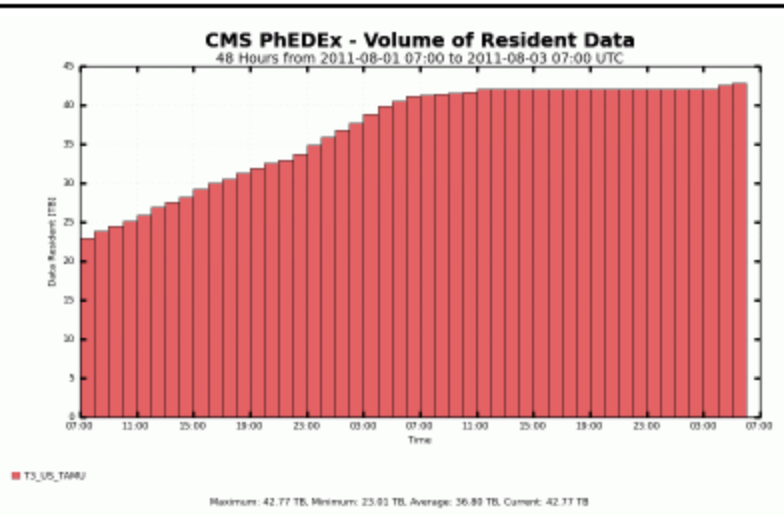	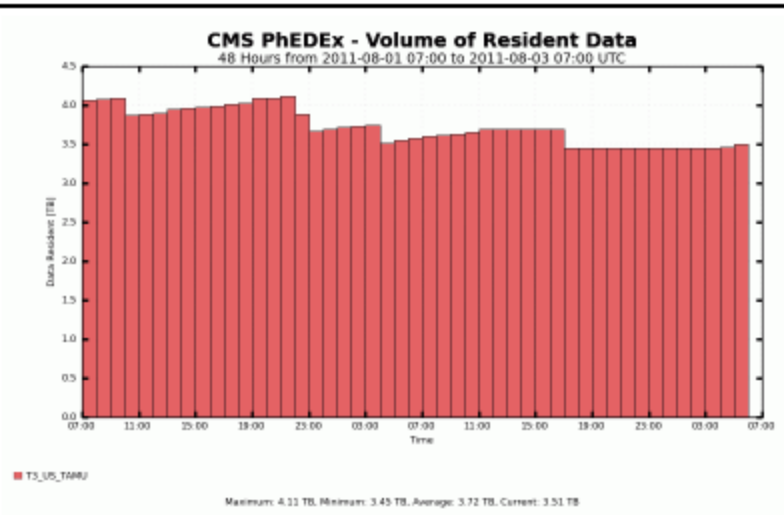
↑ Click to Enlarge Images	Select → Hourly Daily Weekly

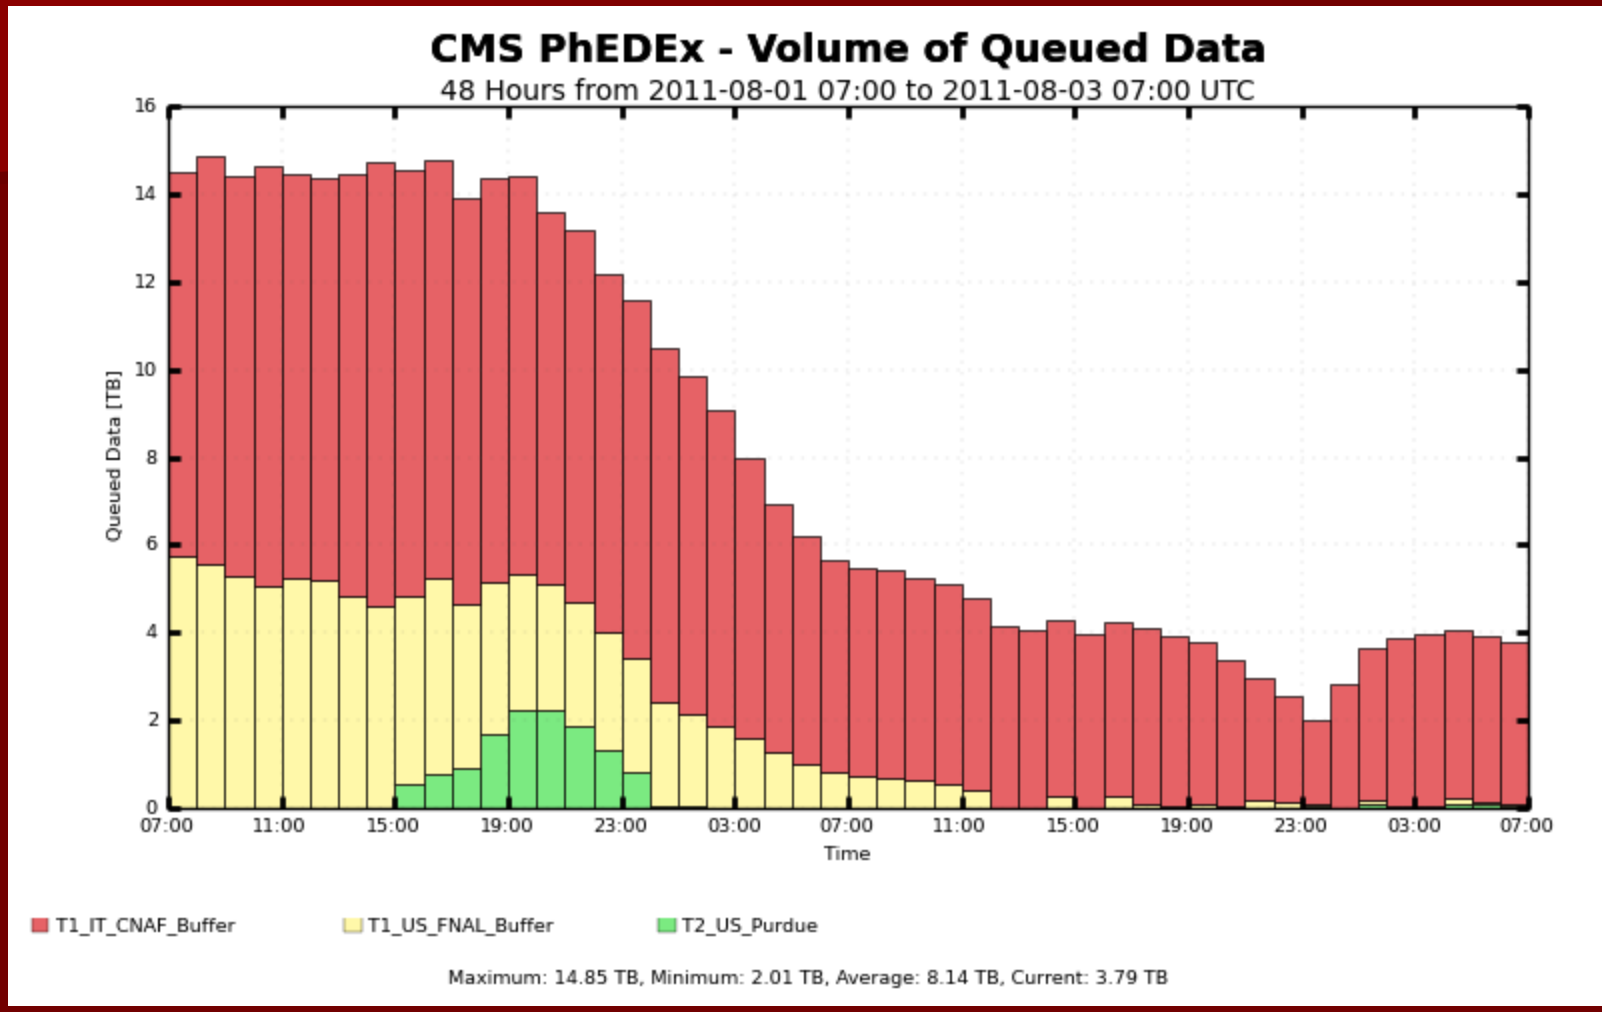
CMS PhEDEx - Volume of Queued Data
48 Hours from 2011-08-01 07:00 to 2011-08-03 07:00 UTC
Queued Data [TB]
0
2
4
6
8
10
12
14
16
07:00
11:00
15:00
19:00
23:00
03:00
07:00
11:00
15:00
19:00
23:00
03:00
07:00
Time
T1_IT_CNAF_Buffer
T1_US_FNAL_Buffer
T2_US_Purdue
Maximum: 14.85 TB, Minimum: 2.01 TB, Average: 8.14 TB, Current: 3.79 TB

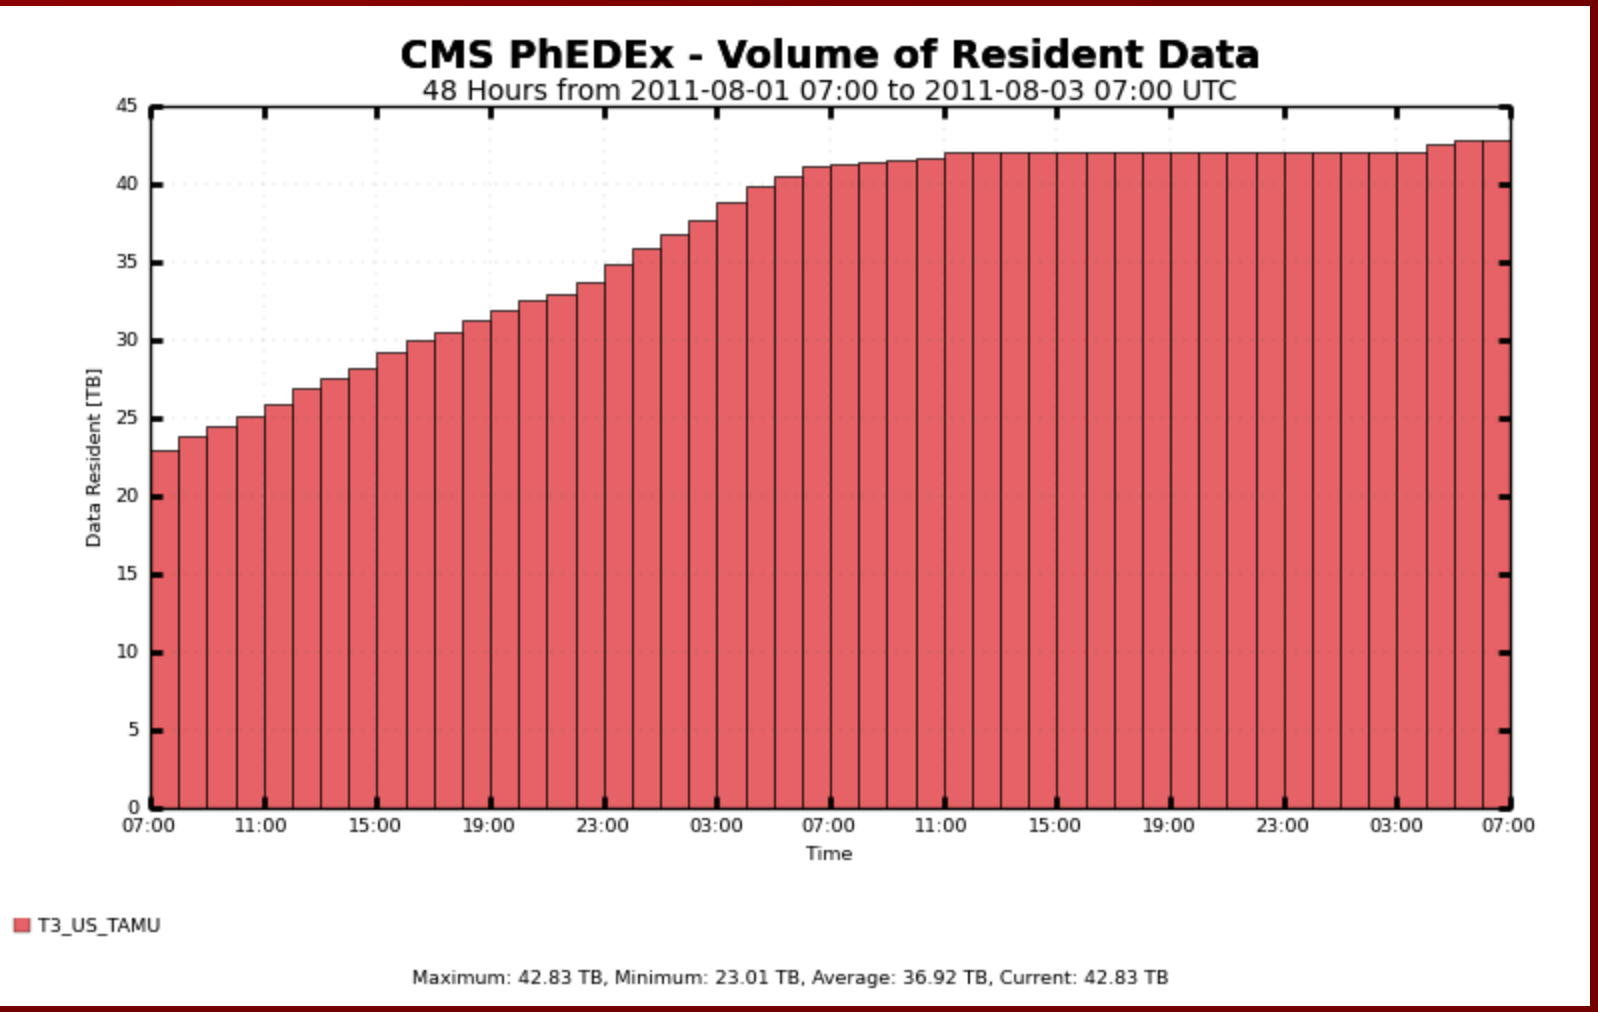
CMS PhEDEx - Volume of Resident Data
48 Hours from 2011-08-01 07:00 to 2011-08-03 07:00 UTC
Data Resident [TB]
0
5
10
15
20
25
30
35
40
45
07:00
11:00
15:00
19:00
23:00
03:00
07:00
11:00
15:00
19:00
23:00
03:00
07:00
Time
T3_US_TAMU
Maximum: 42.83 TB, Minimum: 23.01 TB, Average: 36.92 TB, Current: 42.83 TB

Production Data	Subscribed PhEDEx Transfers			Resident PhEDEx Transfers			
Group Name	Items	Files	Bytes	Items	Files	Bytes	Percent
FacOps	5	1022	2.2 TiB	5	1022	2.2 TiB	100.0 %
deprecated-undefined	1	10	10.2 GiB	1	10	10.2 GiB	100.0 %
exotica	327	7582	16.6 TiB	327	7582	16.6 TiB	100.0 %
susy	7	12189	28.0 TiB	0	10415	23.9 TiB	85.4 %
Totals	340	20803	46.9 TiB	333	19029	42.8 TiB	91.3 %

Data Storage and Monitoring

- Monitor PhEDEx and User files
  - HEPX User Output Files →
  - PhEDEx Dataset Usage ↓

HEPX Disk Store Usage	
Directory	Bytes
User Output	2.6 TiB
Tai Sakuma	1.1 TiB
Indara Suarez	1.0 TiB
Roy Montalvo	436.7 GiB
Alfredo Gurrola	45.4 GiB
Jieun Kim	10.2 GiB
Vaikunth Thukral	20.0 KiB
PhEDEx Monte Carlo	20.3 GiB
PhEDEx Data	20.0 KiB
PhEDEx Load Tests	10.1 MiB
Miscellaneous	404.0 KiB
Total	2.6 TiB
↑ Click to Expand or Collapse Table	

Data on Brazos

	Subscribed			Resident			
Group Name	Items	Files	Bytes	Items	Files	Bytes	Percent
FacOps	1	8	10.1 GiB	1	8	10.1 GiB	100.0 %
deprecated-undefined	2	142	203.3 GiB	2	142	203.3 GiB	100.0 %
exotica	8	5667	8.7 TiB	8	5667	8.7 TiB	100.0 %
higgs	1	10	34.4 GiB	1	10	34.4 GiB	100.0 %
qcd	26	2411	4.5 TiB	16	2348	4.4 TiB	97.4 %
susy	5	6145	9.8 TiB	5	6145	9.7 TiB	100.0 %
Total	43	14383	23.2 TiB	33	14320	23 TiB	99.2 %

Note that this is important for self-imposed quotas. Need to know if we are keeping below our 30TB allocation. Will expand to 50TB soon. Will eventually be sending email if we get near our limit.

HEPX Disk Store Usage	
Directory	Bytes
→ PhEDEx Monte Carlo	26.8 TiB
→ PhEDEx Data	16.0 TiB
→ PhEDEx Load Tests	7.8 GiB
→ User Output	2.7 TiB
→ Miscellaneous	408.0 KiB
Total	45.6 TiB
↑ Click to Expand or Collapse Table	

HEPX Disk Store Usage	
Directory	Bytes
↓ PhEDEx Monte Carlo	26.8 TiB
Summer11	24.6 TiB
Spring11	2.2 TiB
Sam	20.3 GiB
Fall10	200.0 KiB
↓ PhEDEx Data	16.0 TiB
Run2011a	16.0 TiB
→ PhEDEx Load Tests	7.8 GiB
↓ User Output	2.7 TiB
Indara Suarez	1.1 TiB
Tai Sakuma	1.1 TiB
Roy Montalvo	436.7 GiB
Alfredo Gurrola	45.4 GiB
Jieun Kim	10.2 GiB
Vaikunth Thukral	3.7 GiB
↓ Miscellaneous	408.0 KiB
Generators	160.0 KiB
Temp	148.0 KiB
Unmerged	100.0 KiB
Total	45.6 TiB
↑ Click to Expand or Collapse Table	

III – Job Status of the Local Cluster

Process Cycle Statistics | Click to Enlarge Images ↓

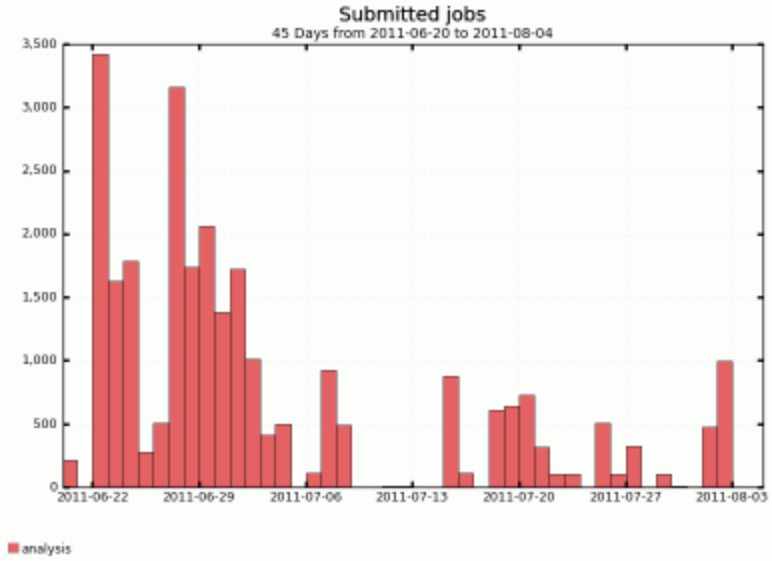

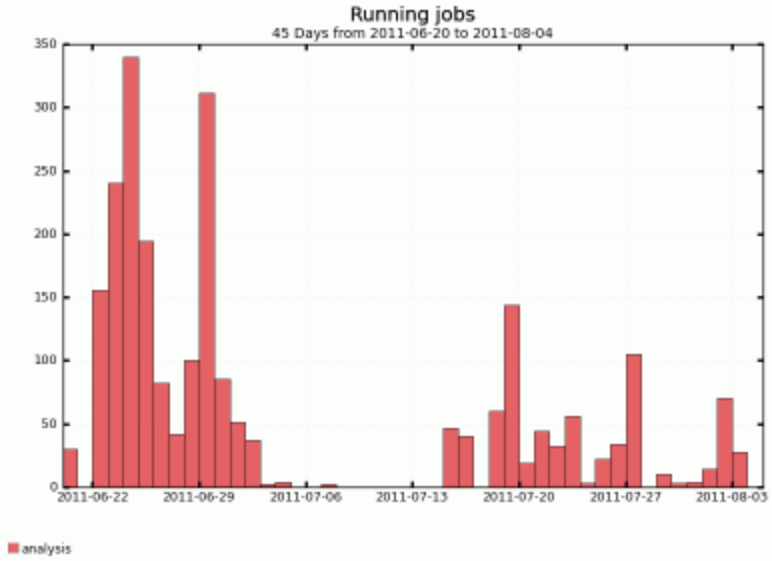

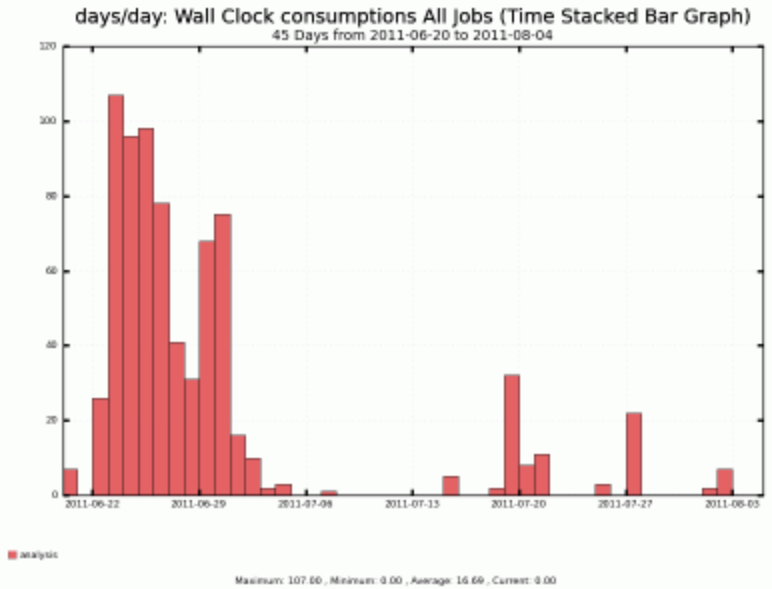

Initiation & Runtime | Terminatation Status | ← Select → | Hourly | Daily | Weekly

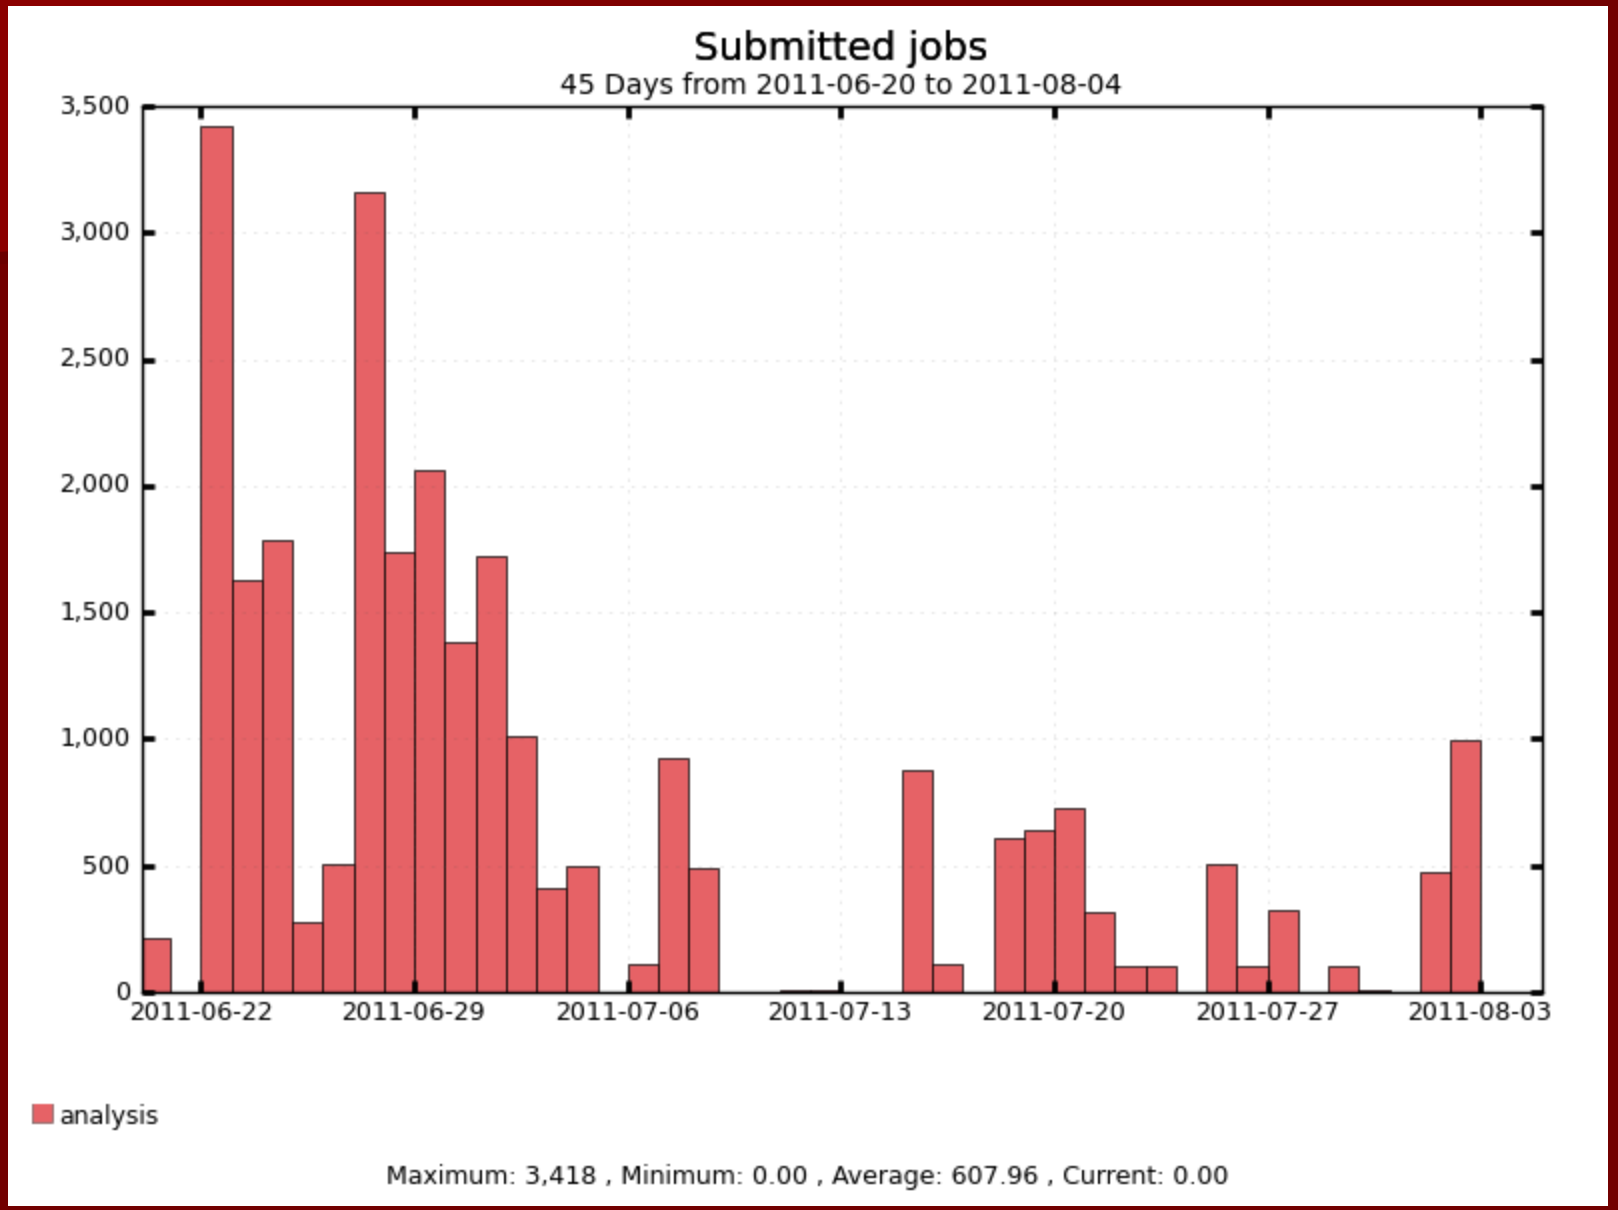
Submitted jobs
45 Days from 2011-06-20 to 2011-08-04
3,500
3,000
2,500
2,000
1,500
1,000
500
0
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
analysis
Maximum: 3,418 , Minimum: 0.00 , Average: 607.96 , Current: 0.00

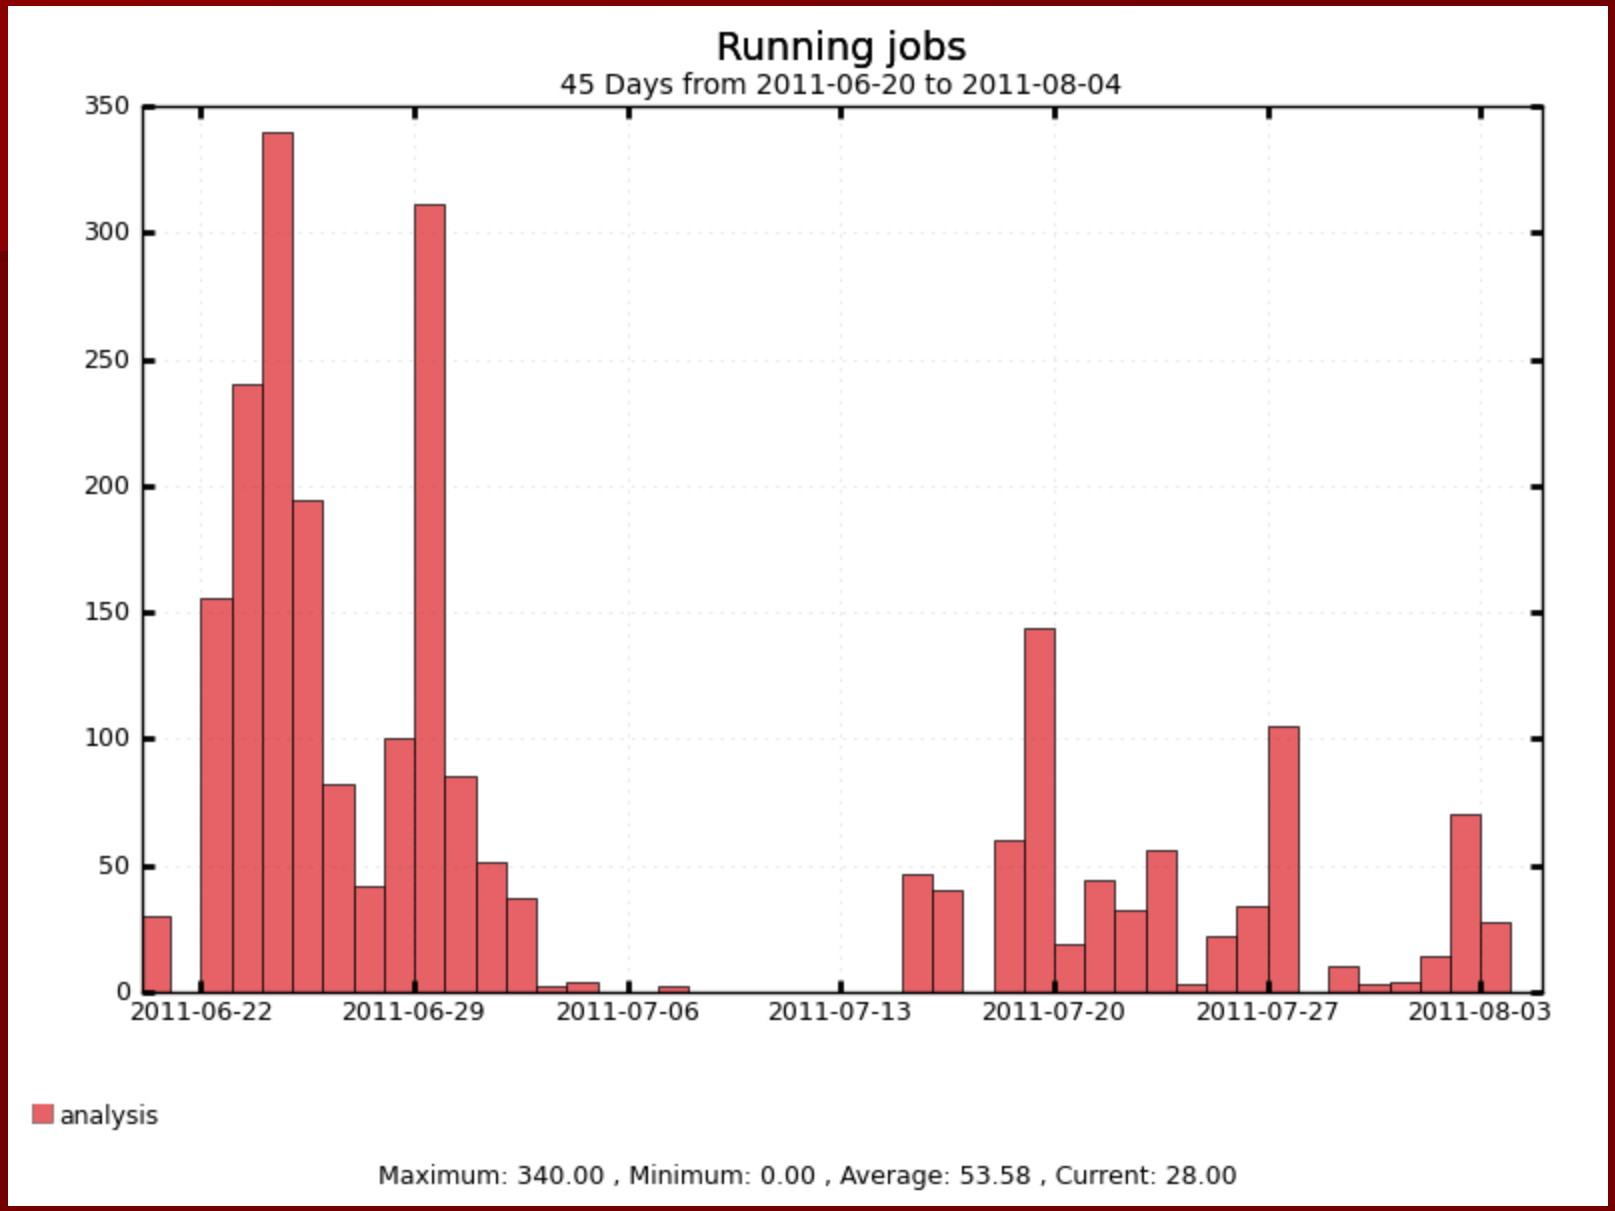
Running jobs
45 Days from 2011-06-20 to 2011-08-04
350
300
250
200
150
100
50
0
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
analysis
Maximum: 340.00 , Minimum: 0.00 , Average: 53.58 , Current: 28.00

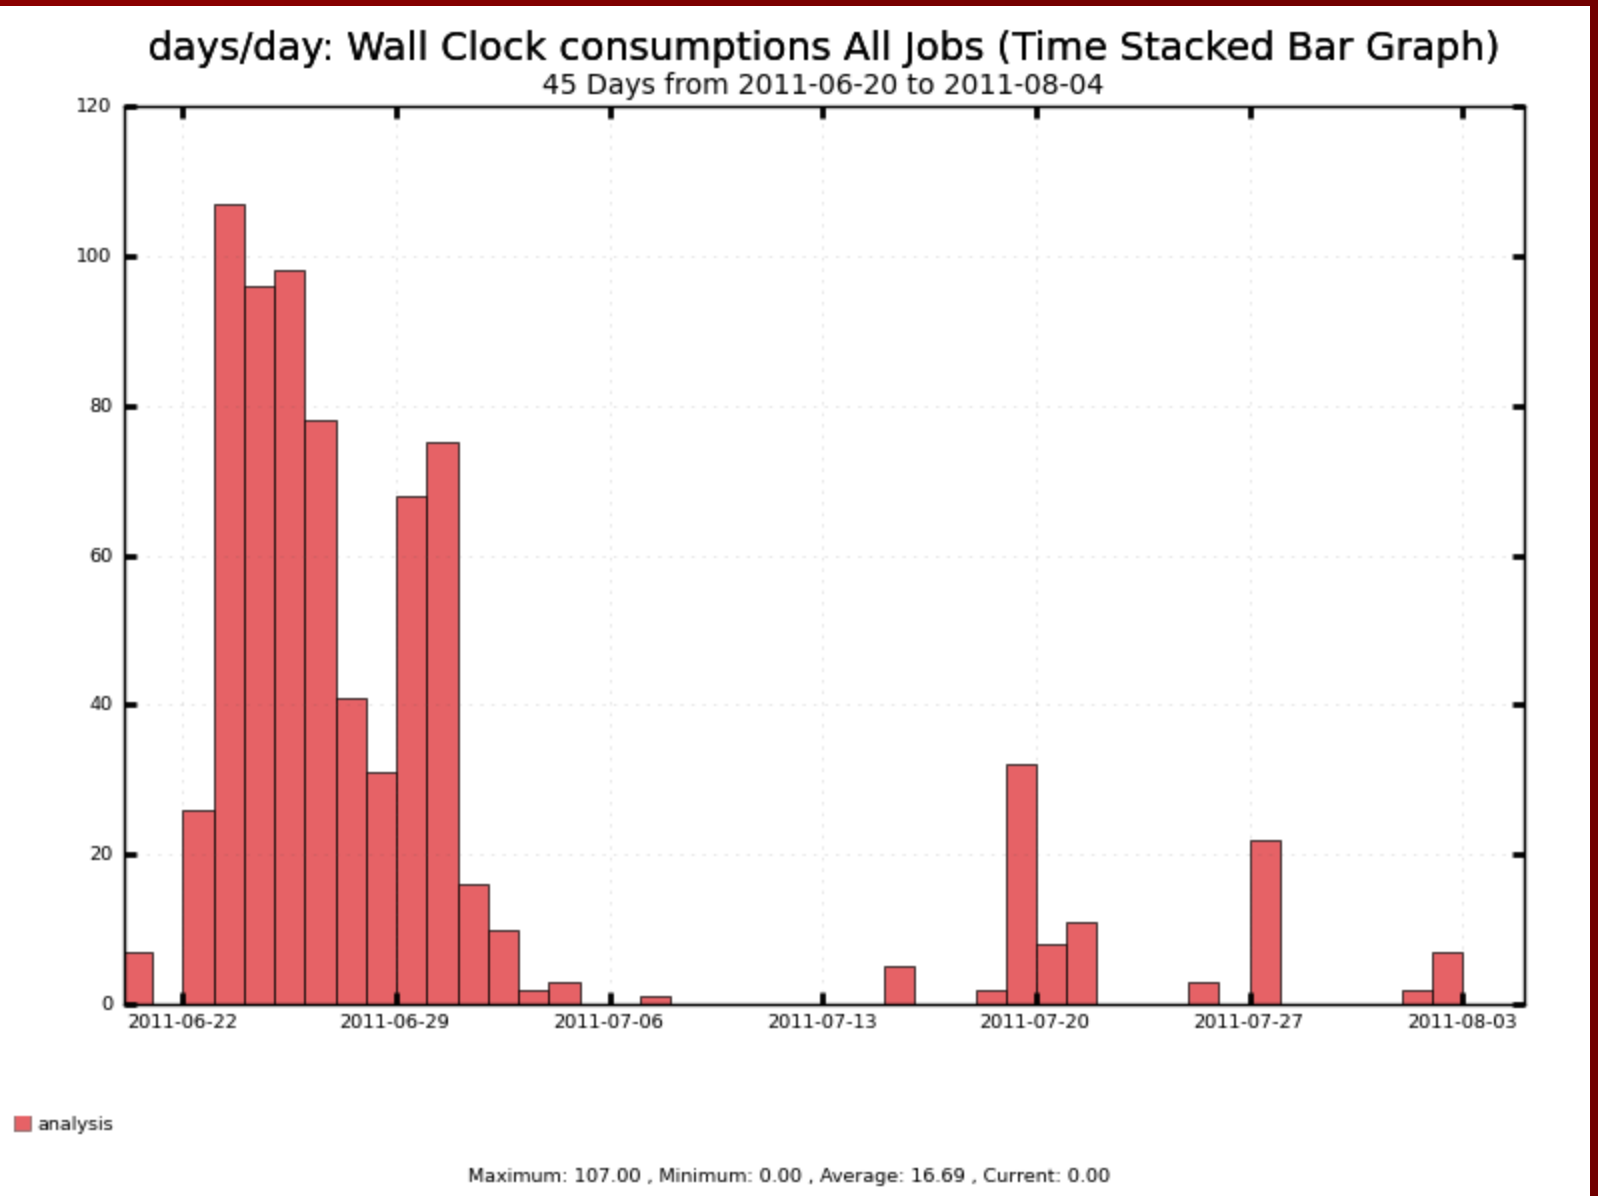
days/day: Wall Clock consumptions All Jobs (Time Stacked Bar Graph)
45 Days from 2011-06-20 to 2011-08-04
120
100
80
60
40
20
0
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
analysis
Maximum: 107.00 , Minimum: 0.00 , Average: 16.69 , Current: 0.00

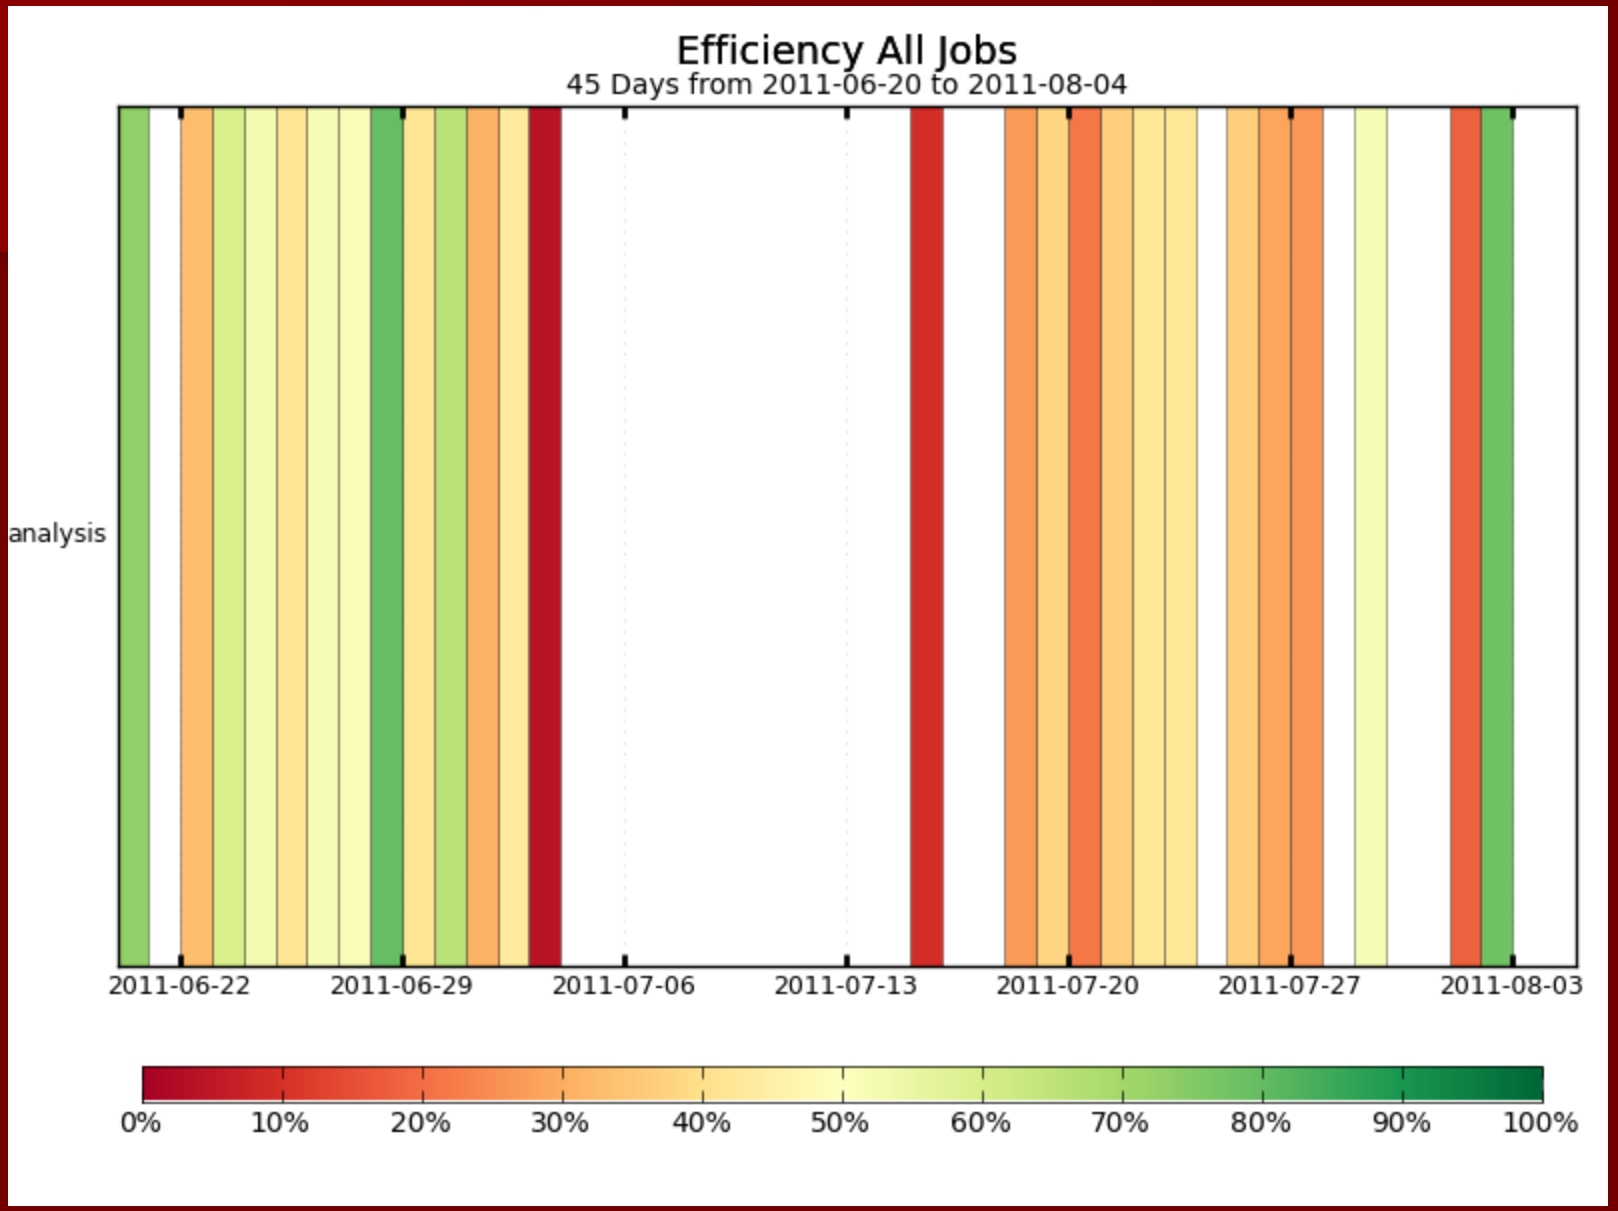
Efficiency All Jobs
45 Days from 2011-06-20 to 2011-08-04
analysis
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%

Process Cycle Statistics

Click to Enlarge Images ↓

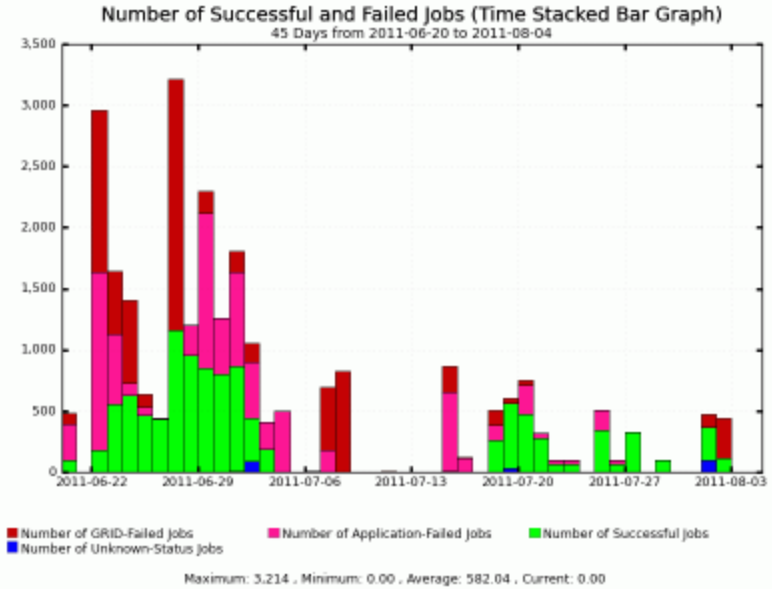

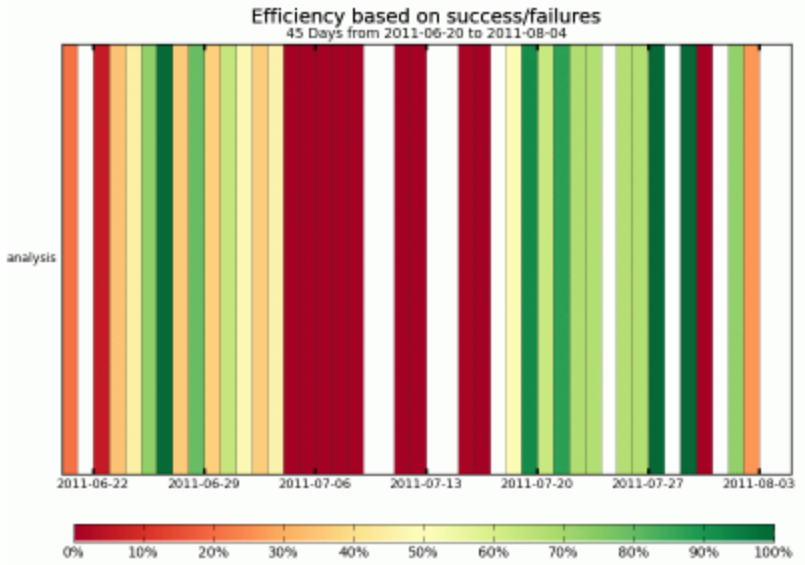

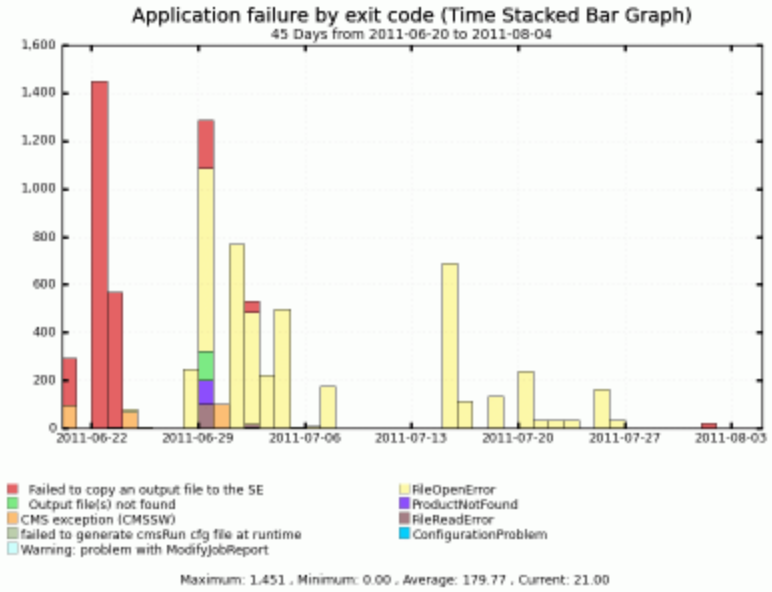

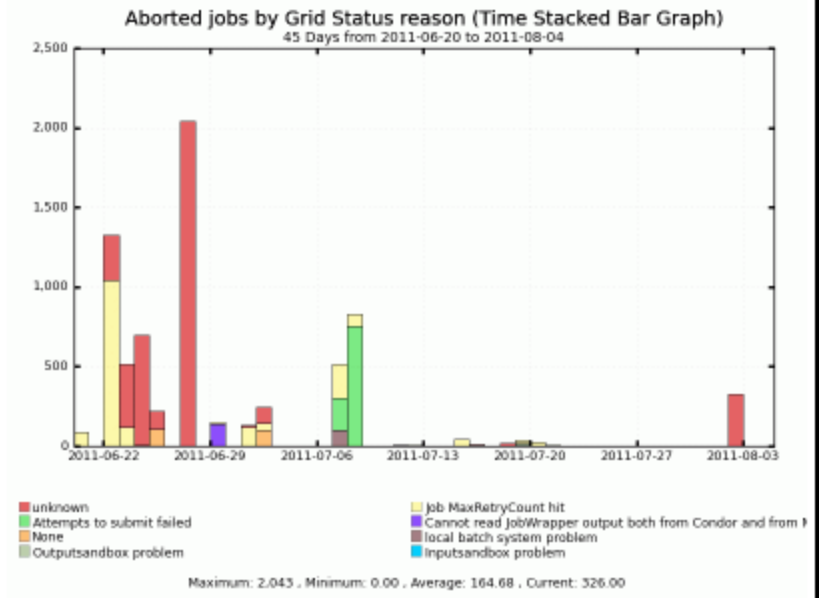

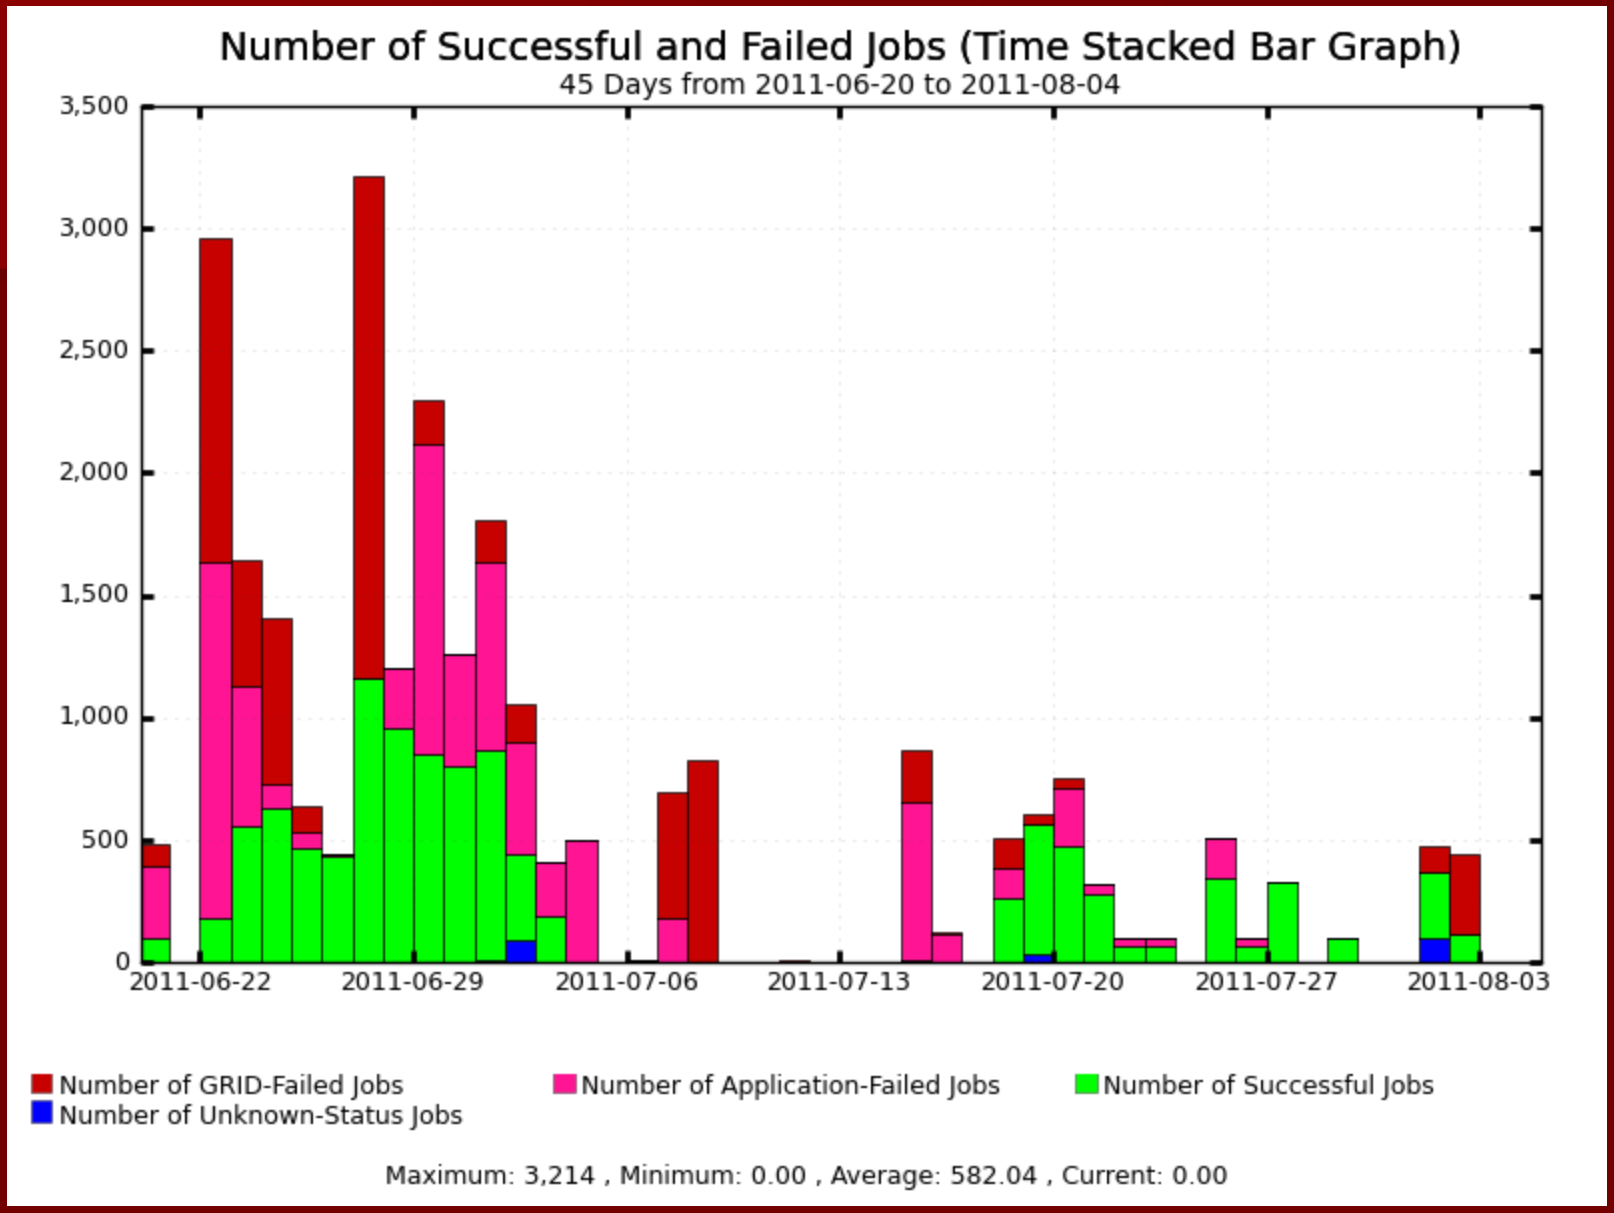
Number of Successful and Failed Jobs (Time Stacked Bar Graph)
45 Days from 2011-06-20 to 2011-08-04
3,500
3,000
2,500
2,000
1,500
1,000
500
0
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
Number of GRID-Failed Jobs
Number of Application-Failed Jobs
Number of Successful Jobs
Number of Unknown-Status Jobs
Maximum: 3,214 , Minimum: 0.00 , Average: 582.04 , Current: 0.00

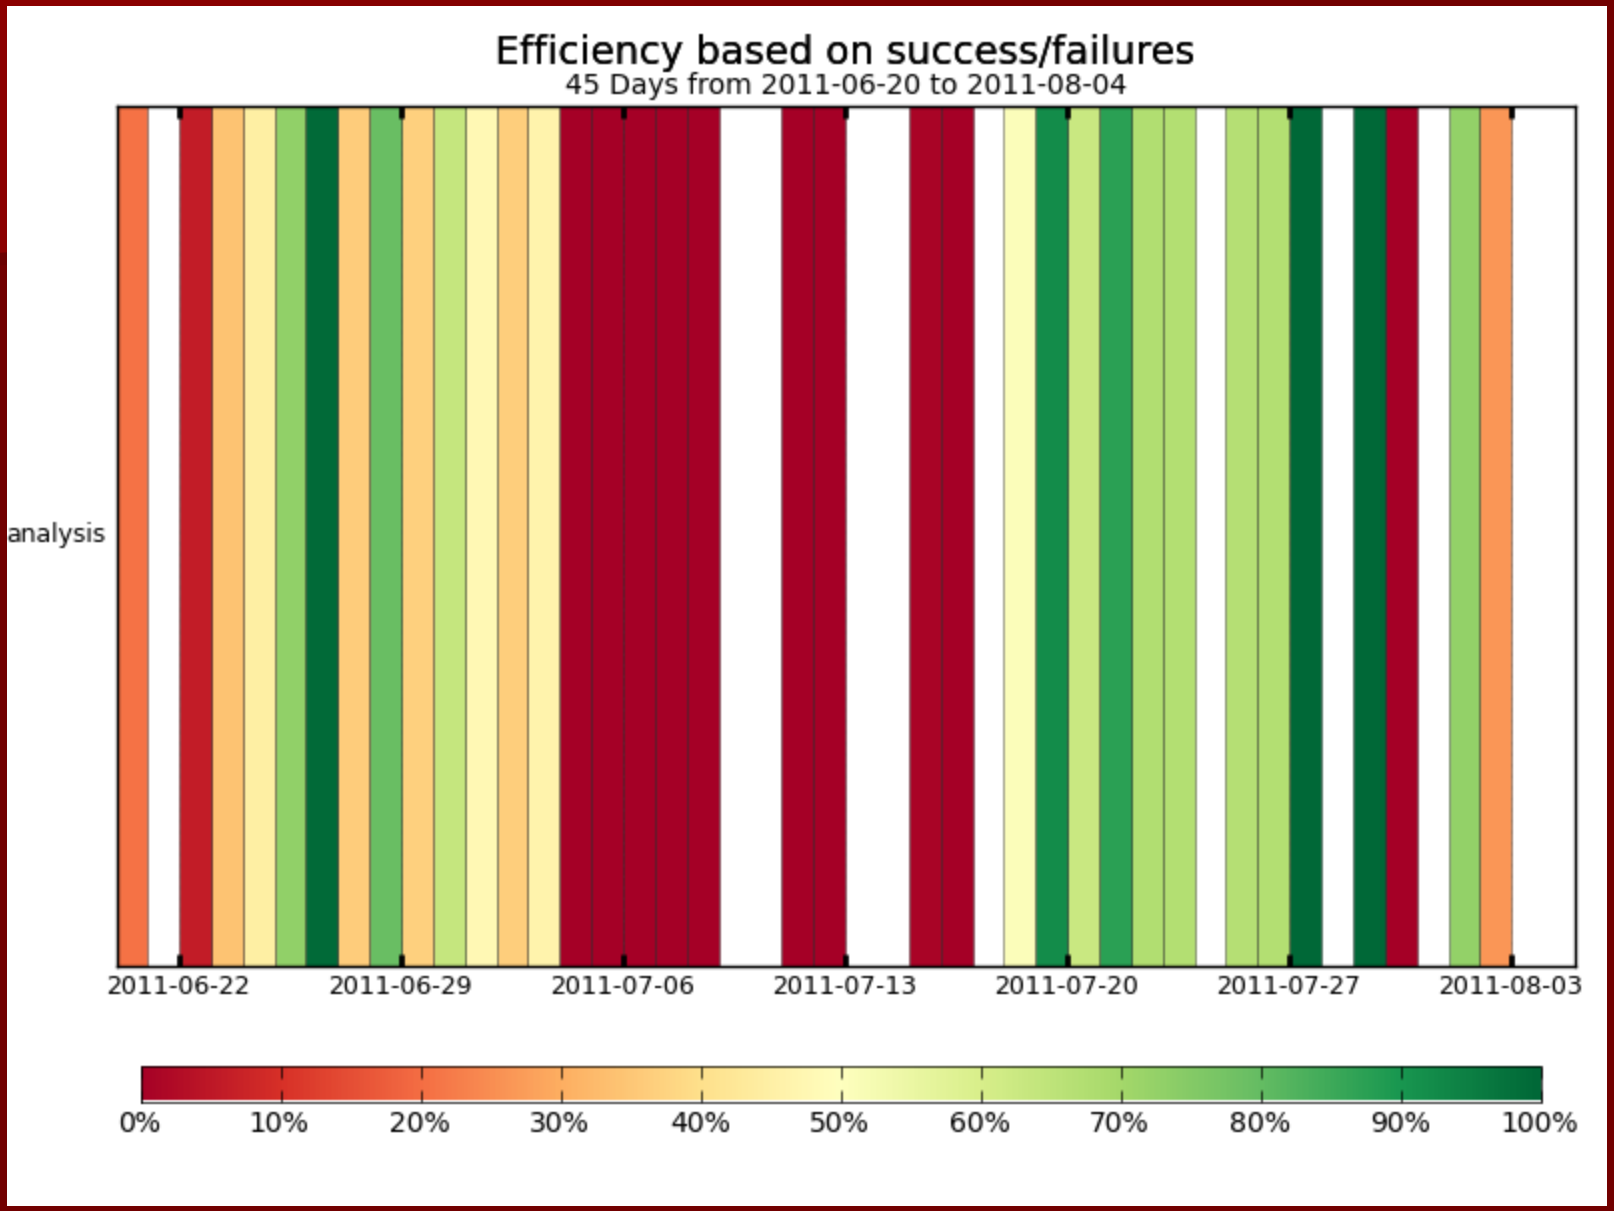
Efficiency based on success/failures
45 Days from 2011-06-20 to 2011-08-04
analysis
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%

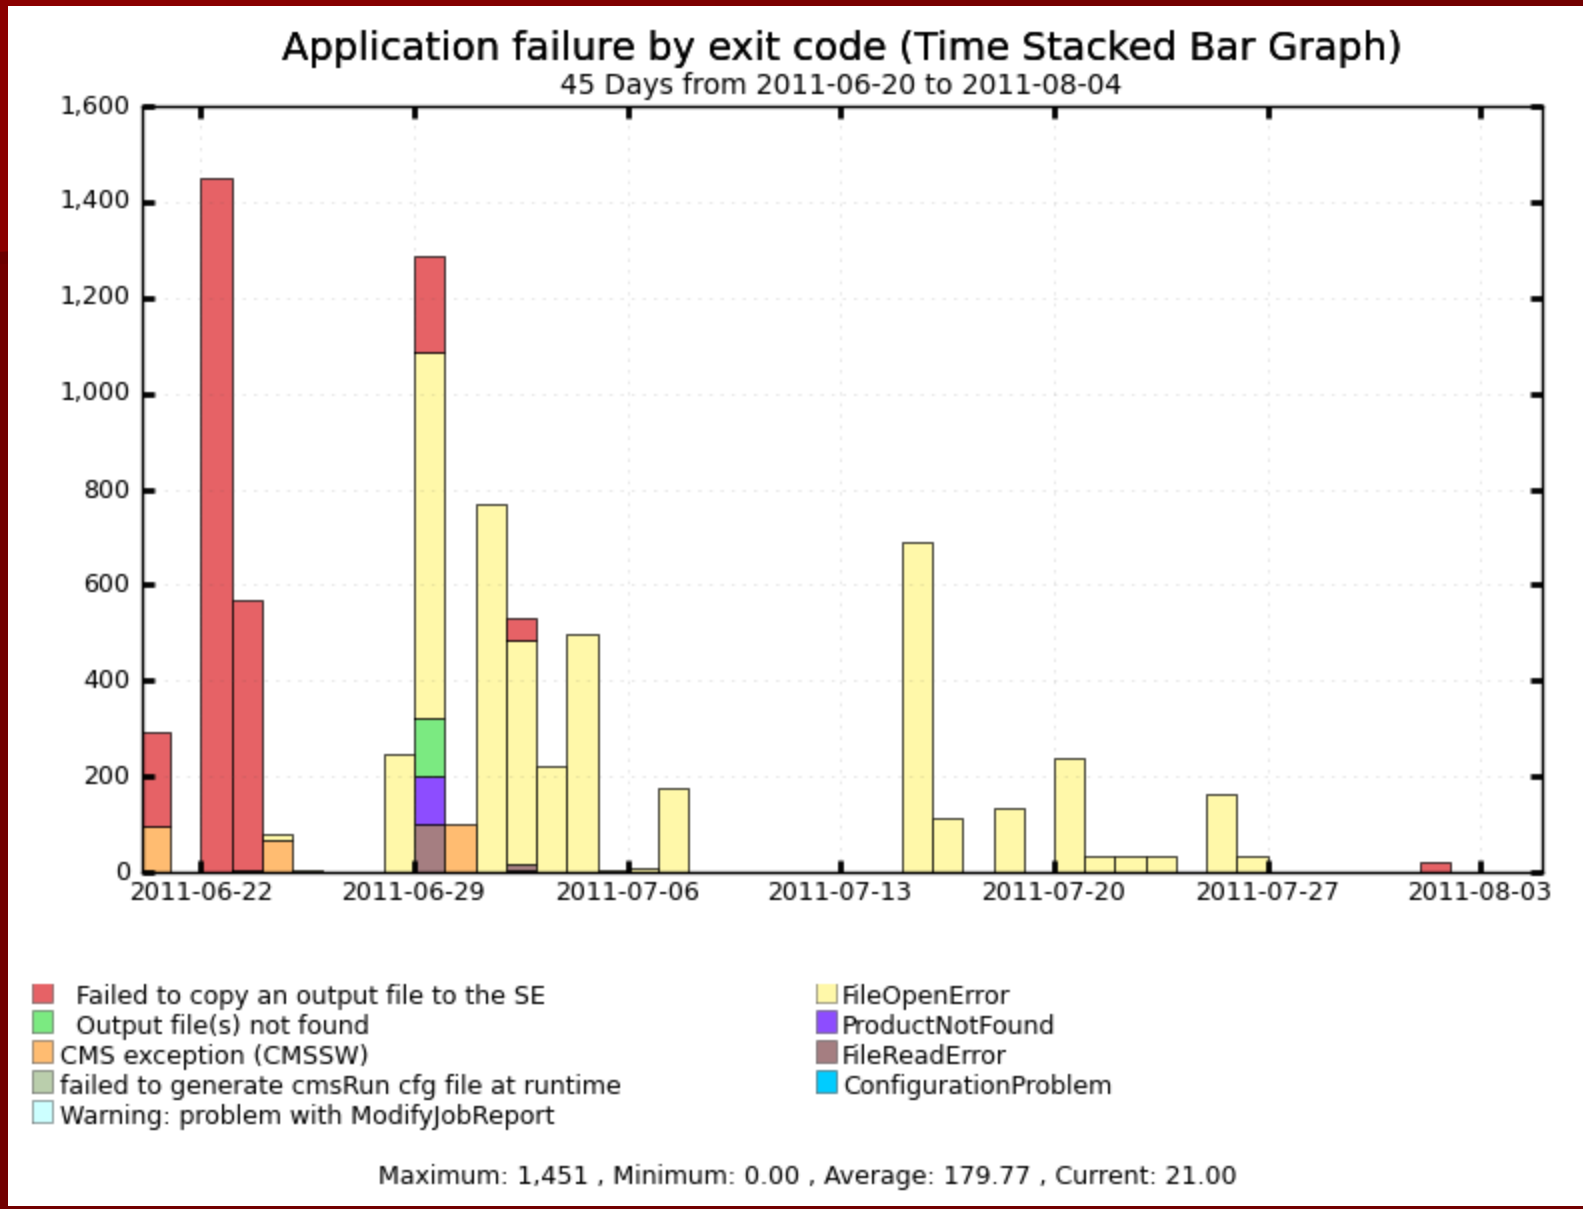
Application failure by exit code (Time Stacked Bar Graph)
45 Days from 2011-06-20 to 2011-08-04
1,600
1,400
1,200
1,000
800
600
400
200
0
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
Failed to copy an output file to the SE
Output file(s) not found
CMS exception (CMSSW)
failed to generate cmsRun cfg file at runtime
Warning: problem with ModifyJobReport
FileOpenError
ProductNotFound
FileReadError
ConfigurationProblem
Maximum: 1,451 , Minimum: 0.00 , Average: 179.77 , Current: 21.00

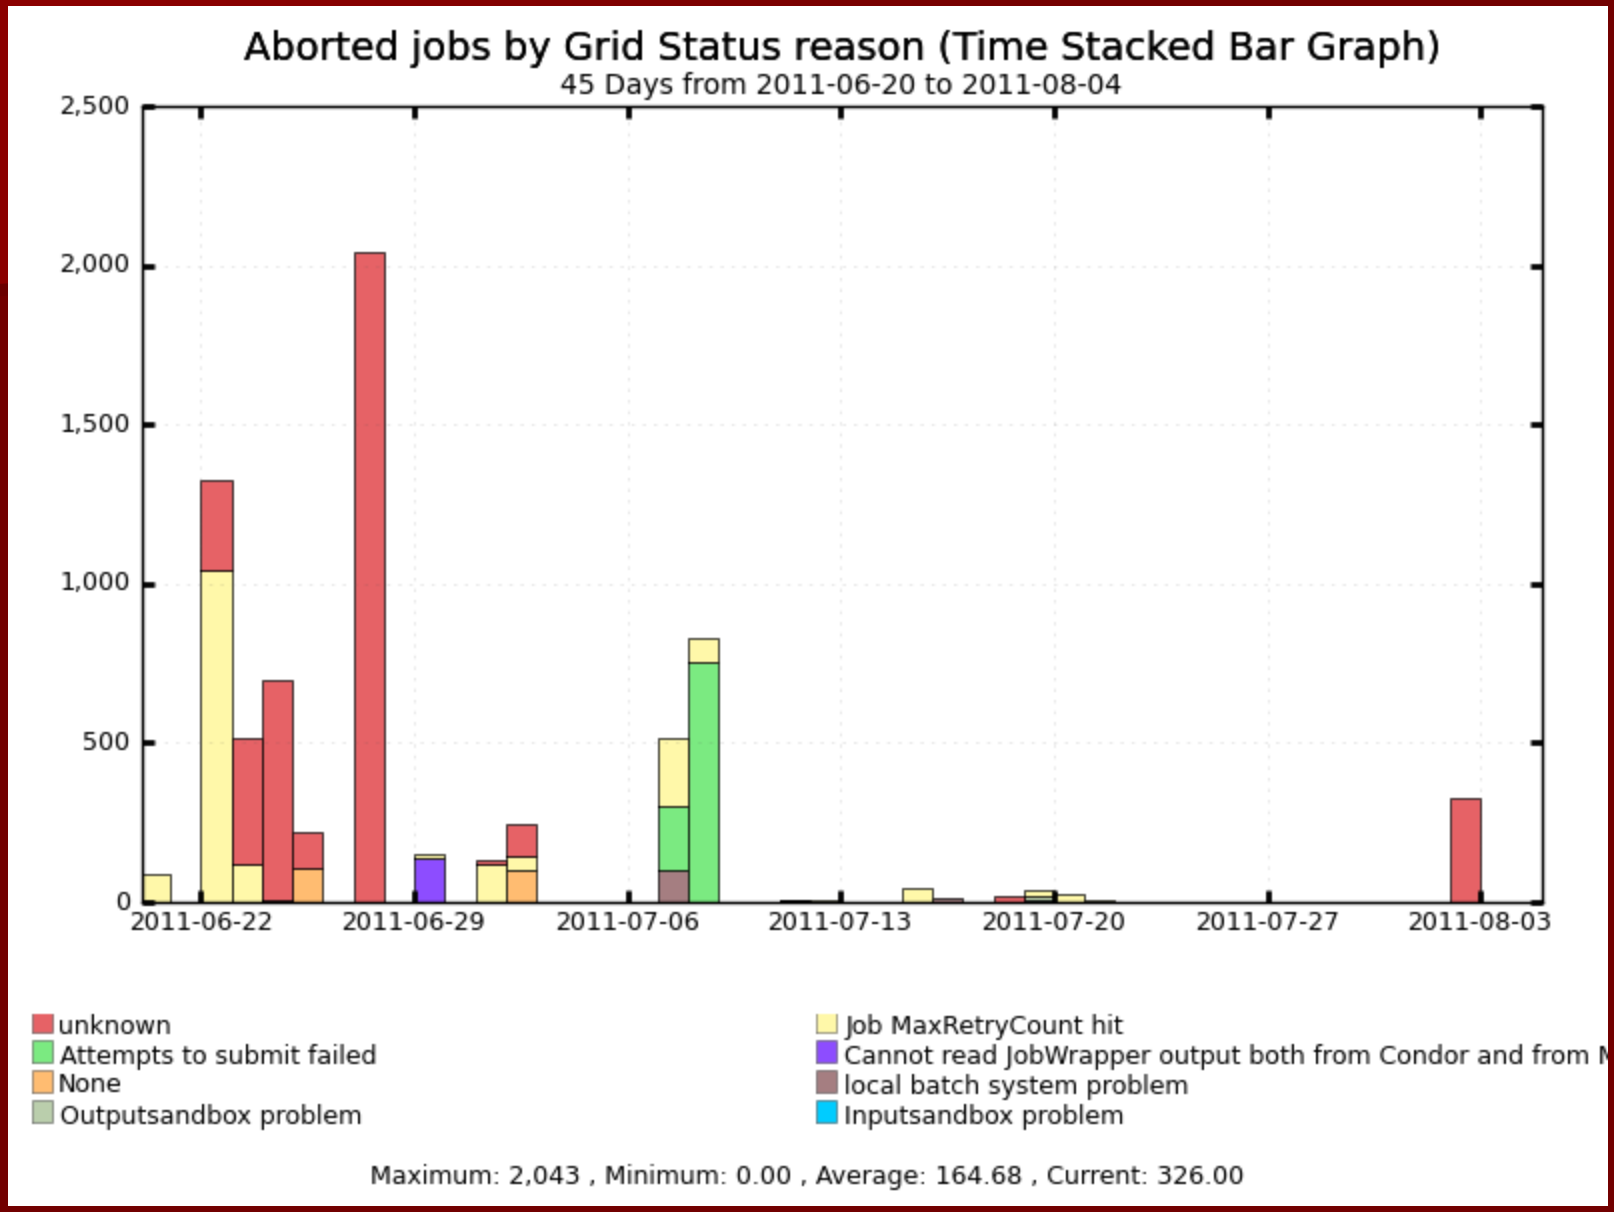
Aborted jobs by Grid Status reason (Time Stacked Bar Graph)
45 Days from 2011-06-20 to 2011-08-04
2,500
2,000
1,500
1,000
500
0
2011-06-22
2011-06-29
2011-07-06
2011-07-13
2011-07-20
2011-07-27
2011-08-03
unknown
Attempts to submit failed
None
Outputsandbox problem
Job MaxRetryCount hit
Cannot read JobWrapper output both from Condor and from M
local batch system problem
Inputsandbox problem
Maximum: 2,043 , Minimum: 0.00 , Average: 164.68 , Current: 326.00

CRAB

- **C**MS **R**emote **A**nalysis **B**uilder
  - Jobs are submitted to "the grid" using CRAB
  - CRAB decides how and where these tasks will run
  - Same tasks can run anywhere the data is located
  - Output can be sent anywhere you have permissions

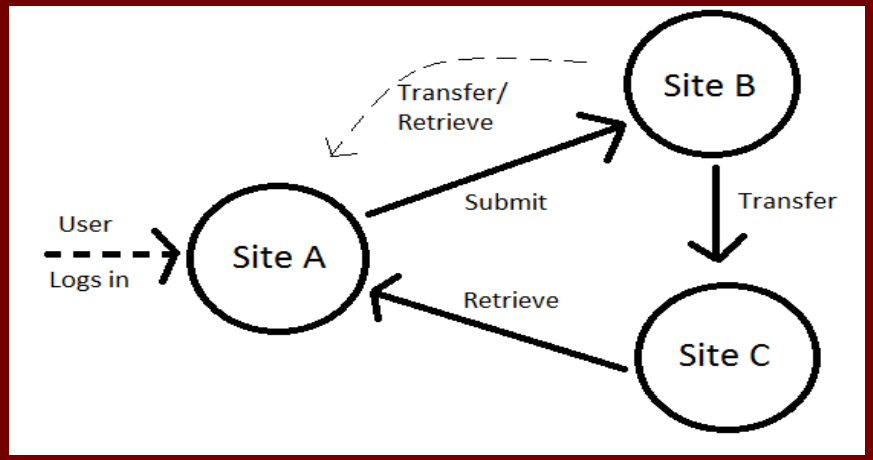

↓ Output Host	Crab Analysis Test Suite (CATS) Completed Job Status			
Grid Client →	gLite		CondorG	
Output Size →	Small	Large	Small	Large
Brazos Cluster	Pass: 5 Fail: 0 Other: 0 2011-08-01 21:33 UTC	Pass: 5 Fail: 0 Other: 0 2011-08-01 21:43 UTC	Pass: 5 Fail: 0 Other: 0 2011-08-03 01:35 UTC	Pass: 5 Fail: 0 Other: 0 2011-08-01 21:33 UTC
Fermi National Laboratory	Pass: 5 Fail: 0 Other: 0 2011-08-01 21:44 UTC	Pass: 5 Fail: 0 Other: 0 2011-08-01 21:42 UTC	Pass: 5 Fail: 0 Other: 0 2011-08-01 21:34 UTC	Pass: 102 Fail: 0 Other: 0 2011-07-27 20:46 UTC
↑ Test Results Link to Job Details ↑				

HEPX & GRID Batch Queue Job Status			
User Name	Jobs	Running	CPU Hours
→ Vaikunth Thukral	5	5	2.4
Totals	5	5	2.4
↑ Click to Expand or Collapse Table			

HEPX & GRID Batch Queue Job Status			
User Name	Jobs	Running	CPU Hours
↓ Vaikunth Thukral	5	5	2.4
HEPX: 815549		Running	0.5
HEPX: 815550		Running	0.5
HEPX: 815551		Running	0.5
HEPX: 815552		Running	0.5
HEPX: 815553		Running	0.5
Totals	5	5	2.4
↑ Click to Expand or Collapse Table			

Service Availability Monitoring (SAM) Tests

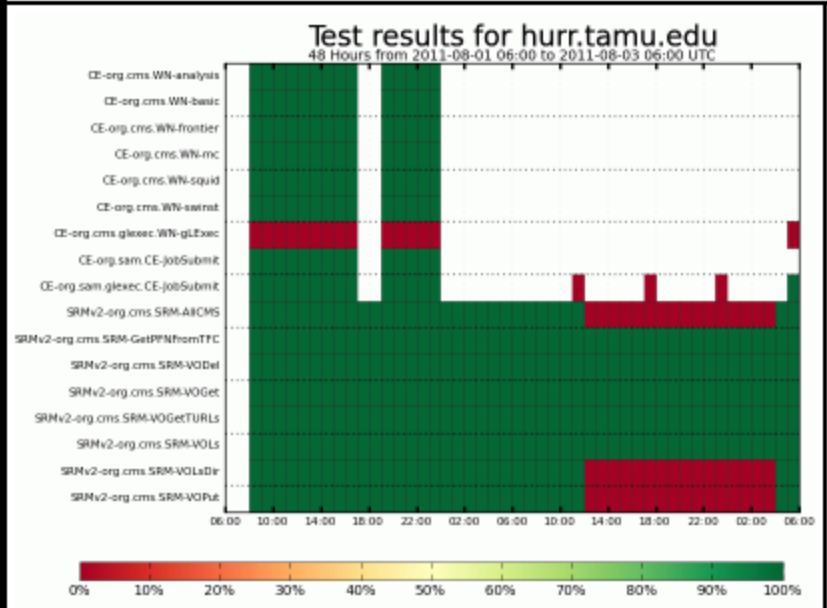

↑ Click to Enlarge Images

Enlarged Images Link to Details ↑

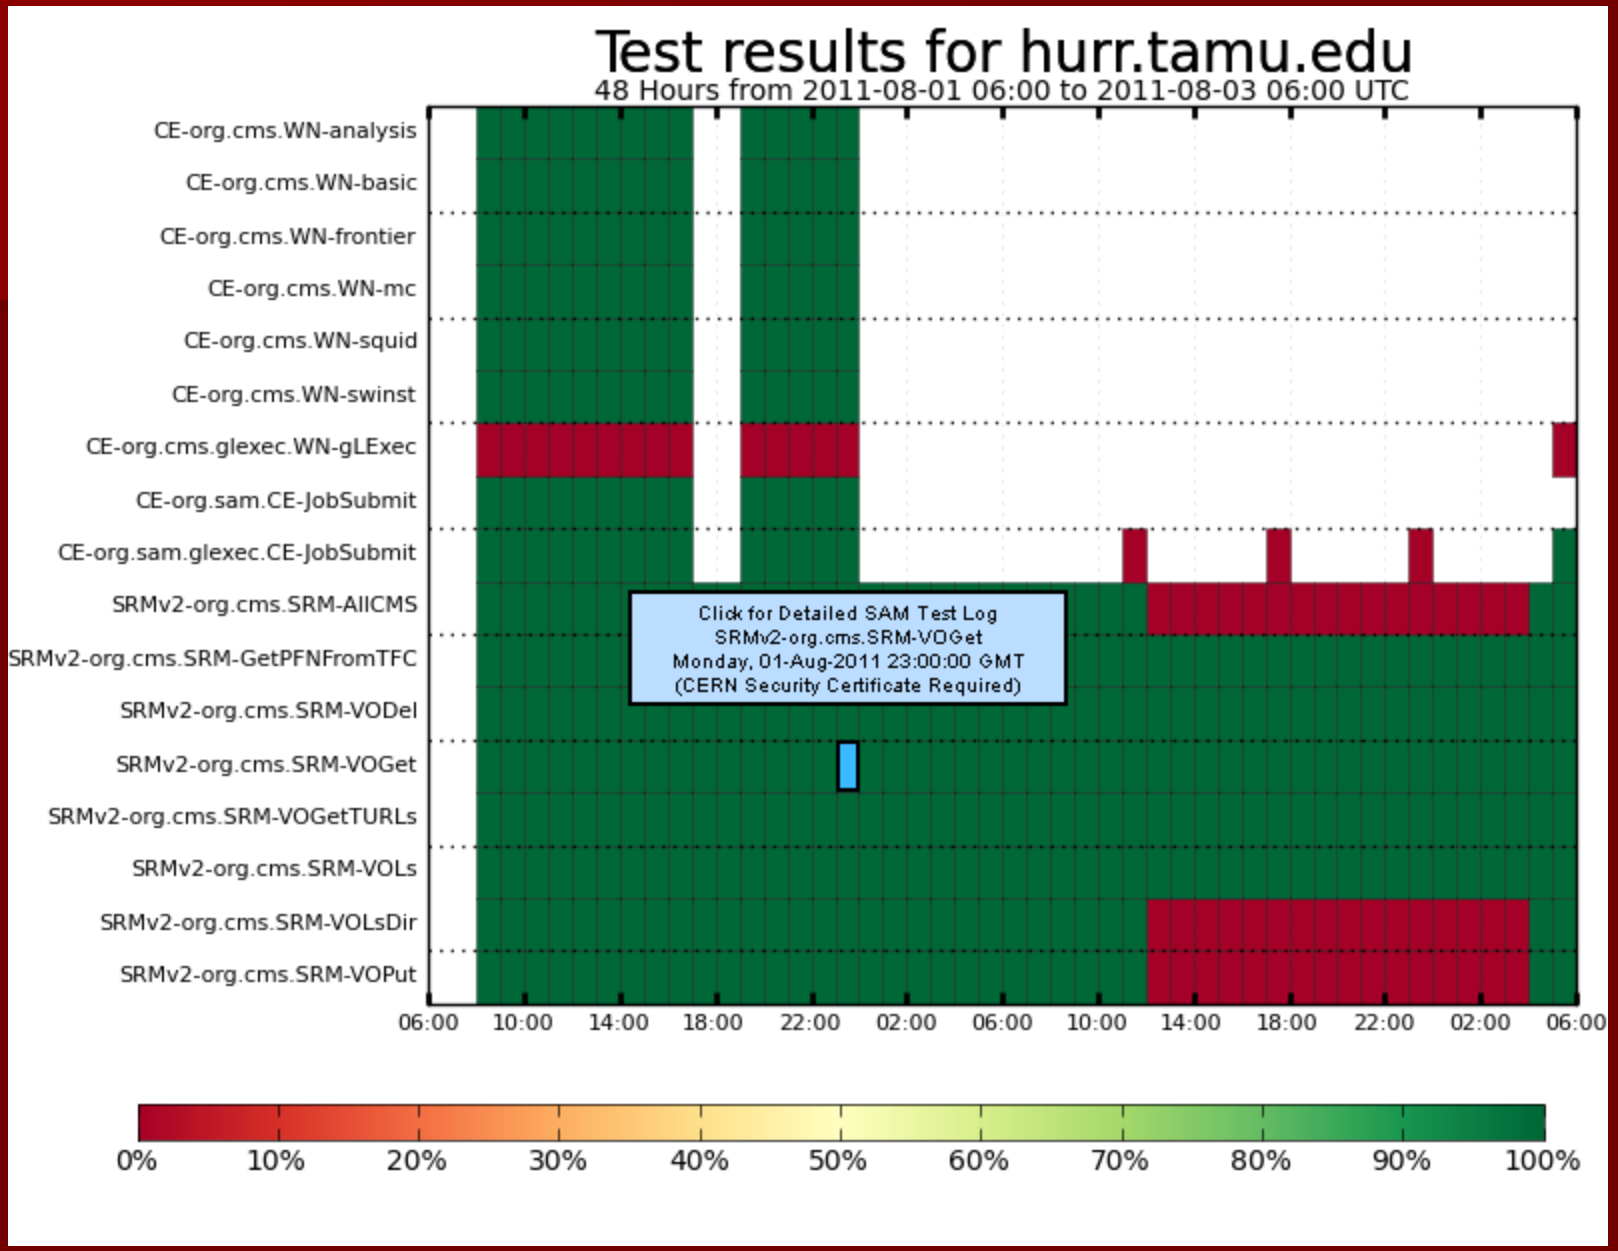
Test results for hurr.tamu.edu
48 Hours from 2011-08-01 06:00 to 2011-08-03 06:00 UTC
CE-org.cms.WN-analysis
CE-org.cms.WN-basic
CE-org.cms.WN-frontier
CE-org.cms.WN-mc
CE-org.cms.WN-squid
CE-org.cms.WN-swinst
CE-org.cms.glexec.WN-gLExec
CE-org.sam.CE-JobSubmit
CE-org.sam.glexec.CE-JobSubmit
SRMv2-org.cms.SRM-AllCMS
SRMv2-org.cms.SRM-GetPFNFromTFC
SRMv2-org.cms.SRM-VODel
SRMv2-org.cms.SRM-VOGet
SRMv2-org.cms.SRM-VOGetTURLs
SRMv2-org.cms.SRM-VOLs
SRMv2-org.cms.SRM-VOLsDir
SRMv2-org.cms.SRM-VOPut
Click for Detailed SAM Test Log
SRMv2-org.cms.SRM-VOGet
Monday, 01-Aug-2011 23:00:00 GMT
(CERN Security Certificate Required)
06:00
10:00
14:00
18:00
22:00
02:00
06:00
10:00
14:00
18:00
22:00
02:00
06:00
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%


Site Availability
45 Days from 2011-06-20 to 2011-08-04
Click for Hourly Summary
All SAM Tests
2011-06-24
T3_US_TAMU
2011-06-23
2011-06-30
2011-07-07
2011-07-14
2011-07-21
2011-07-28
2011-08-
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%

How Much Work Was Involved?

This has been an ongoing project over the course of the Summer of 2011, programmed by myself and my two students Jacob Hill and Michael Kowalczyk, under the close direction of David Toback. Several hundred man-hours have been expended to date.

The critical tasks, above and beyond the physical Perl, JavaScript and HTML coding, include the careful consideration of what information should be included, and how it might most succinctly be organized and presented.

Future Plans

- **Continue to enhance the presentation of our "big three" monitoring targets, and take advantage of the normal "hiccups" in the implementation of a new Tier 3 site to check the robustness and completeness of the monitoring suite.**

- **Implement a coherently managed "Alert Layer" on top of the existing monitoring package.**

- **Seek ongoing funding, and consider the feasibility of sharing the monitoring suite with other Tier 3 sites with similar needs to reduce duplicated workload.**

Alert Layer Design Specifications

- **The alert layer of the Tier 3 monitor system must be organized holistically, and not implemented in a piecemeal fashion spread across various potential alerting tasks.**

- **The alert layer must effectively diagnose “out of bounds” or abnormal behavior, and automatically contact Tier 3 administrators.**

- **The alert layer must be sensitive to context of severity, responding incrementally to increasingly urgent system failures and abnormalities.**

Alert Layer Design Specifications

- **The alert layer must be sensitive to its mailing history, and not “cry wolf” by spamming site administrators with repeated warnings, and particularly not for cases of low severity.**

- **The alert layer must be sensitive to improvement and regression over time, also distinguishing trends with low deviation from stochastic fluctuations.**

- **The alert layer must provide a daily summary of site behavior and alert logging.**

That’s Super, But …

- You need a Snazzy Acronym!

That's Super, But …

- You need a Snazzy Acronym!
- Yes, of course we do! How about the:

"Utility for SAM and PhEDEX Surveillance"

USPS ?

That’s Super, But …

- Can we have it?

That's Super, But …

- Can we have it?

- We would like that, but …
  1. The programs are still in development.
  2. We are not currently funded beyond this summer, and the logistical implications of supporting a network of users need to be considered carefully.
  3. If you like it, you should tell Dave! Or, even better, have your site supervisor drop Dave a line, and they can chat in person: toback@tamu.edu

Summary

- **The Brazos Tier 3 monitoring project has been motivated by the desire to construct a unified, automated repository of up-to-date performance statistics and historical performance calibrations of our local CMS Grid member.**

- **A working "Beta" monitor deployment is already dramatically streamlining the job of keeping tabs on our data transfers, data holdings and job queue status. We expect that this will continue to facilitate the rapid diagnosis and correction of emerging problems.**

Monitoring Your CMS Tier 3 Site

Joel W. Walker

Sam Houston State University

OSG and CMS Tier 3 Summer Workshop

Texas Tech University

August 9-11, 2011

Representing:
the Texas A&M Tier 3 CMS Grid Site on the Brazos Cluster

In Collaboration With:
David Toback
Guy Almes
Steve Johnson
Jacob Hill
Michael Kowalczyk
Vaikunth Thukral (With thanks for marked slides)
Daniel Cruz